

**ORIGINAL RESEARCH**

# Face morphing attacks and face image quality: The effect of morphing and the unsupervised attack detection by quality

Biyang Fu<sup>1</sup>  | Naser Damer<sup>2</sup> <sup>1</sup>Fraunhofer Institute for Computer Graphics  
Research IGD, Darmstadt, Germany<sup>2</sup>Technische Universität Darmstadt, Darmstadt,  
Germany**Correspondence**Biyang Fu, Fraunhofer Institute for Computer  
Graphics Research IGD, Darmstadt, Germany.  
Email: [biyang.fu@igd.fraunhofer.de](mailto:biyang.fu@igd.fraunhofer.de)**Funding information**Hessisches Ministerium für Wissenschaft und Kunst;  
Bundesministerium für Bildung und Forschung**Abstract**

Morphing attacks are a form of presentation attacks that gathered increasing attention in recent years. A morphed image can be successfully verified to multiple identities. This operation, therefore, poses serious security issues related to the ability of a travel or identity document to be verified to belong to multiple persons. Previous studies touched on the issue of the quality of morphing attack images, however, with the main goal of quantitatively proofing the realistic appearance of the produced morphing attacks. The authors theorise that the morphing processes might have an effect on both, the perceptual image quality and the image utility in face recognition (FR) when compared to bona fide samples. Towards investigating this theory, this work provides an extensive analysis of the effect of morphing on face image quality, including both general image quality measures and face image utility measures. This analysis is not limited to a single morphing technique but rather looks at six different morphing techniques and five different data sources using ten different quality measures. This analysis reveals consistent separability between the quality scores of morphing attack and bona fide samples measured by certain quality measures. The authors' study goes further to build on this effect and investigate the possibility of performing unsupervised morphing attack detection (MAD) based on quality scores. The authors' study looks into intra- and inter-dataset detectability to evaluate the generalisability of such a detection concept on different morphing techniques and bona fide sources. The authors' final results point out that a set of quality measures, such as MagFace and CNNIQA, can be used to perform unsupervised and generalised MAD with a correct classification accuracy of over 70%.

## 1 | INTRODUCTION

The advances in the accuracy of FR, driven by innovative training strategies [1, 2] and network architectures [3, 4], are making FR a method of choice for physical and logical access control. However, FR is still vulnerable to attacks such as the face morphing attack. Face morphing incorporates two or more faces from different individuals to create a new face image, such that the newly created face image (the morphed image) can be successfully verified to multiple identities [5]. Therefore, this operation poses a high potential risk in areas such as border control or financial transactions. Typically, a morphed face image does not only has to be verifiable to multiple identities but also has to appear realistic in case of human inspection [6, 7].

Therefore, different studies presented new morphing methodologies that focus on image appearance [6–9].

As the morphing process can leave morphing artefacts, whether performed on the image level or on a representation level in a generative framework, these artefacts can theoretically affect the perceptual quality of the image or even the utility of the image for face recognition. Few previous studies looked at some aspects of the morphed image quality, especially in relation to the bona fide images, but with very limited quality metrics and for the sole reason of proving that the presented morphing processes produce images that are similar to the bona fide ones [6, 7, 10]. Our recent work [11], which this manuscript extends on, went into more details in measuring the effect of morphing on face image quality (FIQ).

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2022 The Authors. *IET Biometrics* published by John Wiley & Sons Ltd on behalf of The Institution of Engineering and Technology.

The work pointed out that the morphing process did result in a consistent effect on certain quality measures, even when re-digitisation is applied to the image. However, this initial work was limited by investigating the effect on only a single landmark-based morphing technique, namely the OpenCV-based morphing [12]. The work did not also consider the possibility of detecting morphing attacks based on quality.

Following the highly interesting outcome of Ref. [11], but the limited generalisability and practicality of its final conclusions, this paper is presented as an extended version of the work done in Ref. [11], which achieved the best poster award at BIOSIG 2021. In this work, we extend the previous contributions as follows:

- Instead of a single morphing technique, this work extends the study of the effect of morphing on face image quality to six different morphing techniques.
- Instead of a single source of bona fide and morphed face images, this work extends that to five diverse morphing datasets, each with multiple variations in the creation or post-processing of the attacks.
- A major extension is the novel exploration of using the quality scores as an unsupervised tool for the detection of face morphing attack and its generalisability across datasets and attack types.

Towards the added value of this extended study, the two main contributions of this work are:

- Theorise a link between face morphing and face image quality and utility. Based on that, we uncover and analyse the effect of face morphing on face image quality and utility and the stability of this effect across different morphing methods.
- Propose the use of face image quality and utility measures as an unsupervised MAD method by taking advantage of the morphing process effect on the quality measures. We additionally investigate the generalisability of this proposed MAD concept across various variations that might face MAD in real applications.

These contributions thus drive two basic research questions addressed in this work: (1) what is the effect of morphing on face image quality and utility? and (2) can this effect be leveraged as an indication to perform MAD in an unsupervised manner? To address these questions and reach to a clear confirmation on our contributions, we perform experiments on 10 quality measures, divided into face image quality assessment (FIQA) and general image quality assessment (IQA) categories, experimented on 5 morphing datasets that include six different morphing techniques and five different data sources. Our analyses towards answering the main research questions involved the following main detailed investigations:

- Studying the effect of a wide set of morphing techniques on the image qualities, measured by a diverse set of quality measures, in comparison to different bona fide sources.

- Revealing the different levels of separability between quality scores produced by different quality measures on bona fide versus morphing attacks.
- Investigating the quality differences between different bona fide sources when the quality is measured by different quality measures.
- Probing the possibility of leveraging the quality effect to detect face morphing attacks (i.e. classifying bona fide and morphing attacks) by using the quality scores as an MAD decision score within each morphing dataset.
- Investigating the generalisability of classifying bona fide samples of different sources when the quality scores are used as an MAD decision score and decision thresholds are set on other unknown bona fide sources.
- Investigating the generalisability of detecting morphing attacks from different morphing techniques and unknown data sources when the decision thresholds are set on other setup containing unknown morphing attacks and bona fide sources.
- Revealing the overall possible performance and generalisability of MAD based on different quality metrics on a wide set of attacks and data sources.

Our study points out that for certain quality and utility metrics, such as CNNIQA [13] and MagFace [14], there is a significant effect of morphing on the estimated qualities. Such methods also produced similar quality values for different sources of bona fide images. Using selected quality measures can lead to overall MAD accuracy above 70% on a diverse set of unknown attack methods and data sources, without the explicit MAD training.

We structured our work in the following by first introducing relevant studies in Section 2 with respect to studies relating image quality to morphing attacks, where the quality of morphs are however mostly compared to bona fide source images and the detection of morphing attacks. In Section 3, we describe the quality and utility measures studied in this work to investigate how the morphed images affect the image quality and if these findings can be leveraged on the MAD task. The morphing datasets and their used morphing techniques are presented in Section 4 followed by the experimental design with the used evaluation metrics and the investigation protocols. Finally, the results are thoroughly discussed in Section 5 and concluded in the Section 7.

## 1.1 | Nomenclature

To enable a clear understanding of the experimental design and discussion, we introduce here a list of the used terminologies that are essential to be differentiated.

- A morphing ‘dataset’ refers to a complete set of data that includes both the morphing attacks and the bona fide images. The used morphing datasets are detailed in Section 4.2.

- A ‘data source’ is the face image dataset where the bona fide samples (and the samples used to create the morphing attacks) of each morphing dataset originate.
- A ‘morphing technique’ is a process that creates a morphing attack from two or more bona fide samples. The morphing methods considered in this work are presented in details in Section 4.1.
- A ‘morphing attack’ is an image resulting from morphing images belonging to multiple identities and aim at being positively verifiable by these identities.
- An ‘image quality assessment method (IQA)’ is a method that assesses the perceived or statistical quality of an image and not specifically designed to measure the utility of the image for face recognition.
- A ‘face image quality assessment method (FIQA)’ is a method that assesses the utility of a face image for face recognition.
- An ‘unknown morphing method’ is a morphing method not used in the training of a specific MAD and thus ‘unknown’ to the MAD approach.
- An ‘unknown dataset’ is a morphing dataset that is not used in the training of a specific MAD and thus ‘unknown’ to the MAD approach.
- An ‘unknown data source’ is a data source that is not used in the training of a specific MAD and thus ‘unknown’ to the MAD approach.
- The ‘intra-dataset detectability’ is the experimental setup where the performance of MADs is measured on the same morph dataset used for training, although on an identity-disjoint test set not used for training.
- The ‘inter-dataset detectability’ is the experimental setup where the performance of MADs is measured on morphing datasets not used for training the MADs.

## 2 | RELATED STUDIES

Face recognition systems (FRS) are highly vulnerable to face morphing attacks [15]. This poses a severe security risk to operations using automatic face recognition systems for authentication. Therefore, it is of great interest to understand the effect of such attacks and to have effective defence mechanisms. This section of related work thus gives a brief but concise view on two aspects, the previous studies touching on the relationship between morphing and image quality and a high-level view on MAD.

### 2.1 | Morphing attacks and quality

In Ref. [10], Debiasi et al. investigated the qualities on morphing attacks, by using image quality (IQ) measures such as BIQI [16], BRISQUE [17], OG-IQA [18] and SSEQ [19]. These methods are all based on studying the general image qualities and belong to the group of no-reference blind IQA methods, which require no bona fide source reference image as a comparison to assess the morphing attack’s image quality.

Debiasi uses these quality measures to investigate the vulnerability of the morphing attacks on the face recognition solutions. The intuition behind this is based on Ferrara et al.’s work [15] pointing out that human experts have difficulties recognising morphed face images of high quality and bona fide images. Therefore, the authors herein [10] believed that the quality of morphs might be linked to image quality. The experiments are conducted on the MorGAN dataset, which contains the bona fide images, MorGAN, and OpenCV generated morphs. Results show similar quality distributions of MorGAN attacks to the bona fide images within the MorGAN dataset, while landmark-based attacks are slightly off in terms of image quality. This makes the detectability of these GAN-based attacks more difficult compared to classical approaches.

Regarding generating morphs with high quality and resolution, to make the morphing attack more effective, the work in Ref. [6] proposed an approach to enhance the GANs generated images into a more realistic quality and resolution morphing attacks. This improvement could further suppress the artefacts but not affect their performance on the face recognition system. To evaluate the quality of the generated attacks, single quality measures are used as suggested in Ref. [20]. The authors reported multiple quality factors, such as sharpness [21], blur [22], exposure [23], global contrast factor (GCF) [24], contrast [21, 25], and brightness [25]. These quality measures are reported both for bona fide and morphing attacks. The intuitions behind these quality factors are, the closer the measure to the bona fide quality factor, the better the generated morph image. Based on these outcomes, Ref. [6] showed that the images enhanced with the proposed method are of higher visual quality compared to unprocessed attacks and attacks enhanced by other state-of-the-art super-resolution methods.

Zhang et al. [7] proposed a new morph generation approach using an Identity Prior Driven Generative Adversarial Network (MIPGAN) to create high quality and high resolution morphed facial images with minimal artefacts. This is enabled by using a specific loss function, enforcing high perceptual quality and identity factors to keep generated morphs without loss of identity. The MIPGAN Face Morph Dataset is created from FRGC-V2 face dataset using only high-quality passport-like face images from 140 data subjects. The analysis of the perceptual image quality of morphs is performed by comparing it to the reference image. This is contradictory to the other previously mentioned quality measures, which favoured the no-referenced quality assessment methods. The peak signal-to-noise ratio (PSNR [26]) and the structural similarity index measure (SSIM [27]) are used to evaluate the quality of the generated attacks. Results based on these quality measures showed (1) little deviation in the perceived quality for all four different types of morphing attacks and (2) the MIPGAN created morphs have slightly higher image quality compared to StyleGAN 2, and (3) MIPGAN morphs present similar quality to facial landmark-based morphing attacks. In addition, the authors asked 56 human observers to make a visual comparison between pairs of images to assess the morphing qualities besides the quantitative measures such as PSNR and SSIM. They found out that the

detection of morphed face images is a challenging task even for human observers, both the experienced and inexperienced user groups when showing only one single image at a time.

The recent work [11] investigated in total 12 quality measures on the LMA-DRD morphing dataset shows the effect of the morphs on quality and utility. The morphing dataset contains only one morphing technique based on OpenCV. To measure the face utility, four FIQA methods (MagFace [14], SER-FIQ [28], FaceQnet [29], and RankIQ [30]) are utilised to determine a quality value, which is directly related to the term face utility for an FR algorithm. Overall, eight general IQA methods are used to assess the perceived image quality on morphing attacks. These IQA methods include BRISQUE [17], PIQE [31], NIQE [32], CNNIQA [13], DeepIQA [33], MEON [34], RankIQA [35], and dipIQ [36]. All selected quality and utility measures do not require a bona fide source to evaluate the attack's quality. Their results stated that most investigated IQ measures show only minor shift in quality for bona fide images compared to morphing attacks. Only MagFace [14] from the FIQA methods shows a clear separability between morphing attacks and bona fide images. Less differentiation between the re-digitisation process and the digital image is detected by MagFace in terms of separability, which points out that the difference is correlated to the morphing process itself, while most IQA solutions do differentiate between the digital and re-digitised images but less on morphing attack itself. However, the results are limited to a single morphing technique and bona fide source. No studies on the detectability of morphing attack based on these quality and utility measures were conducted.

While most previous studies used quality measures to support their claim of producing attacks similar to bona fide images, the contributions in Ref. [11] that we substantially extend here were limited to studying the effect on one morphing technique and one data source, without considering the generalisability of this effect, nor leveraging this effect to be used as an unsupervised MAD methods. This motivates our work to present novel contributions in these regards.

## 2.2 | Morphing attack detection

MAD methods can be categorised into two main groups, single image and differential MAD [37]. Single image MAD only analyses the investigated image to make a predicted decision of attack or bona fide [38–42]. Differential MAD analyses an investigated image along with a live image (assuming that the process allows for that). Differential MAD analyses the relation between both images to build a decision of attack or bona fide [43–46]. However, the applicability of differential MAD is limited by the requirement of a live image and thus might be less practical in various use-cases.

Single image MAD solutions are commonly modelled as a binary classifier and can be roughly categorised into ones using handcrafted features and ones using deep learning features. Such handcrafted features included Binarised Statistical Image Features (BSIF) [47, 48], Local Binary Patterns (LBP) [49],

Local Phase Quantisation (LPQ) [41], or features established in the image forensic analyses such as the photo response non-uniformity (PRNU) [10]. The MAD solutions based on deep learning commonly used pre-trained networks with or without fine-tuning, such as versions of VGG [50], AlexNet [42], or networks trained for face recognition purposes such as OpenFace [40]. However, all these studies used a single binary label as the target of their training. A recent work leveraged the use of pixel-wise supervision towards creating a more generalisable MAD [39]. The work showed higher generalisability over re-digitised samples when using the proposed training paradigm. The use of synthetic-based face images to create a diverse MAD development data was recently proposed and proved to enhance the performance and generalisability of MADs in comparison to MADs developed based on authentic data [51], especially given that the authentic data is constrained by its lack of diversity and the ethical/privacy issues related to using and sharing it.

A number of MAD studies have analysed the issue of the generalisability of the MAD decisions when facing variabilities in the face morphing or image handling process. Such variabilities included the synthetic image generation processes [6–8, 52], different data sources [53], morphing pair selection [44], image compression [54], and re-digitisation [42, 47, 50]. These variabilities have been shown to cause a drop in the MAD performance when they were unknown in the MAD training phase. A research direction towards avoiding the dependency on training data, and thus enhance the generalisability in detecting attacks with unknown variabilities, is the unsupervised detection of attacks as anomalies. MAD based on anomaly detection was only rarely addressed in the literature [52] based on one-class classification of handcrafted and deep learnt features, however, with very limited detection performance [52]. This work, in its effort to investigate the possibility of using FIQA and IQA as indicators of an image being morphed, also offers an unsupervised MAD method that operates without the need to train an MAD to detect already known attacks.

## 3 | FACE IMAGE QUALITY AND GENERAL IMAGE QUALITY

FIQ represents the utility (value) of the face image to FR algorithms [55, 56]. This utility can be measured by the FIQ score (scalar) following the definition in ISO/IEC 2382-37 [56]. The algorithms designed to measure this utility metric belong to the family of FIQA methods. The general image quality, on the other hand, focusses on perceived image quality and it does not necessarily reflect the utility of the image in FR, for example, an image of a mask occluded face can be of high perceived quality but of low utility to FR algorithms [57]. To investigate the effect of face morphing on the image and the possibility to use this effect in detecting morphing attacks, we select a total of 10 different quality measures, divided between IQA and FIQA categories, which will be summarised in this section.

The morphing processes do introduce changes to the bona fide morphed images. These images can be, to variant degrees, apparent according to the morphing process, see Figures 1 and 2. We theorise that the artefacts, whether minor or major, might have an effect on the perceived image quality measured by IQA. These artefacts, as well as the fact that the image becomes less distinct (can be matched to multiple identities), are also theorised to have an effect on the utility of the image in FR, which can be measured by FIQA. This leads to the wide range of investigations in this work, addressing both the effect of different morphing techniques on the IQ and FIQ of the images as well as the possible use of this effect in detecting morphing attacks.

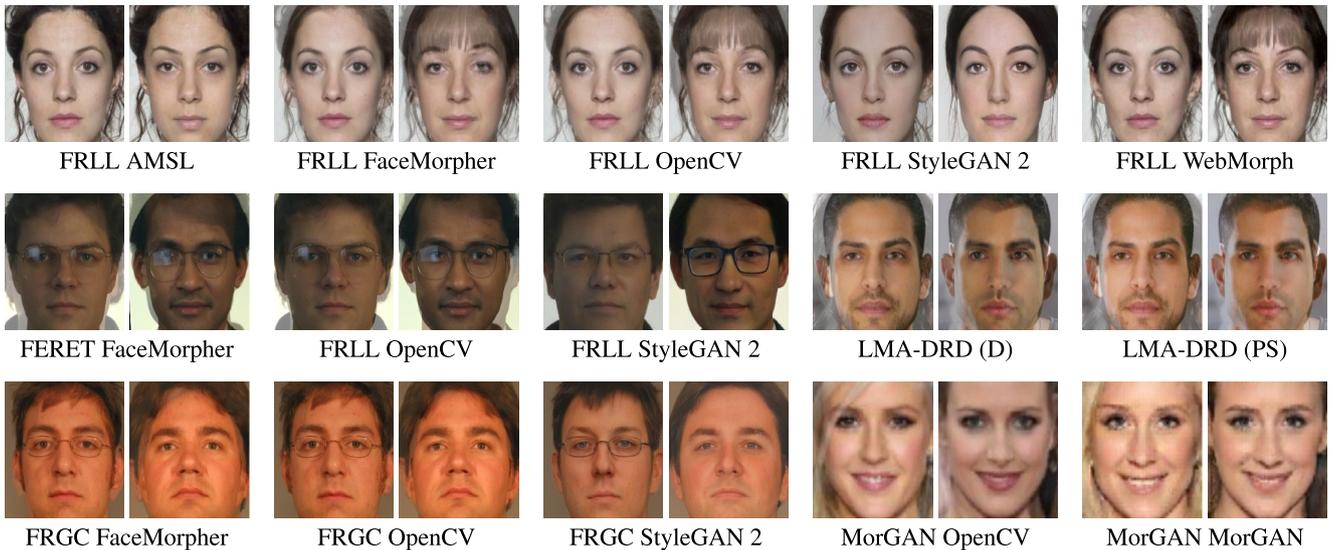
### 3.1 | General image quality assessment

The five IQA approaches selected in this paper can be grouped into the following four sub-categories: (1) based on natural image statistics [17], (2) convolutional neural network-based

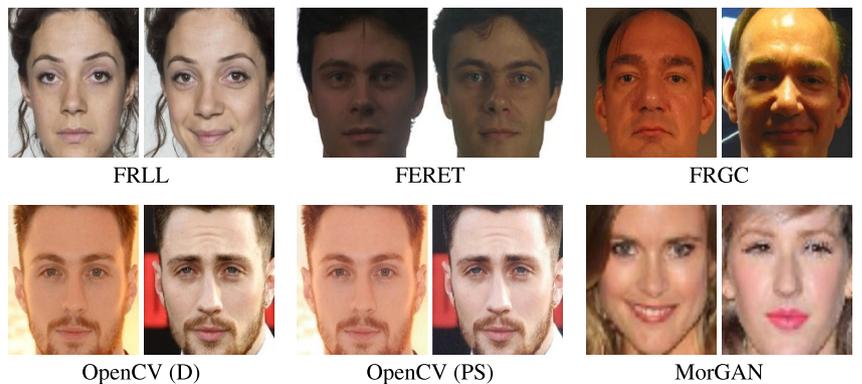
[13], (3) multi-task learning-based approach [58], and (4) ranking-based methods [51, 88]. These assessed qualities by these methods are typically affected by image distortions, artefacts, and perceived degradation.

**BRISQUE** [17] accounts for the first sub-category and is learnt based on studying the deviation from the general statistics of natural images. These statistics are originated from the finding by Rudermann [59] that assumes natural scene images have a distribution similar to a normal Gaussian distribution. The degree of deviation from the normal Gaussian is thus directly related to the degradation in image quality.

The second sub-category uses convolution layers to automatically extract deep features without the need for a priori special design using handcrafted features. **CNNIQA** [13] is accounted to this category. This approach is a patch-based approach that correlates the image quality based on the fusion of patch decisions, enabling local decisions to communicate to form a global decision. The default setting with a patch size of  $32 \times 32$  and a kernel size of  $7 \times 7$  is chosen according to the original publication in Ref. [13].



**FIGURE 1** Each pair of images is an example of the morphed images generated by the morphing attack strategies from the considered 5 morphing datasets. The sub-caption first presents the name of the database followed by the naming of the morphing strategy. Visible artefacts in the morphed images are most apparent for landmark-based morphing techniques especially close to the nose and eye region, while the GAN-based approaches generate morphed images with clearer background



**FIGURE 2** Each pair of images presents the bona fide images from the 5 morphing datasets

However, we applied the sliding window approach with a stride size of 4 pixels to increase the patch resolution as the input image has only a reduced image resolution compared to the training images used in Ref. [13].

**DBCNN** [58] is another general IQA method based on multi-task learning. In the core of the algorithm, two CNNs were trained, one on large-scale synthetically generated datasets, while the other focusses on the classification network pre-trained on ImageNet [60] for extracting more authentic distortions. The final quality score is pooled from the features of both CNNs to a unified representation.

The methods **UNIQUE** [61] and **dipIQ** [36] are categorised into ranking-based IQA approaches. One benefit of these methods is that it is easier to generate ranked image pairs since the absolute quality score for each image is hard to acquire. In addition, it is simple to introduce additional image distortions and synthesise ranked image pairs. Therefore, the cost of generating a large, synthetic dataset with ranked image pairs is relatively low. The network structures of these methods are built using two parallel streams with shared weights. Only one trained stream is used to assess the quality of the input image in the prediction stage.

### 3.2 | Face image quality assessment

Recent FIQA methods take advantage of deep-learning-based approaches and a large amount of publicly available face image datasets. Thus, these methods show relatively good performance in predicting the utility of a face image for FR. The five selected FIQA methods are to be categorised into training with (1) supervised regression learning based on pseudo quality labels as in Ref. [14, 62], (2) unsupervised methods based on FR model behaviour as in Ref. [29, 63], or (3) ranking-based approaches as in Ref. [30].

We used two unsupervised FIQA methods, namely **MagFace** and **SER-FIQ** proposed by Meng et al. in Ref. [14] and Terhoerst et al. in Ref. [62]. In MagFace, the FIQ estimation is considered by the loss function during the training process. The face embedding learnt by the network can be used for both FR and FIQ estimation. The magnitude of the unnormalised face embedding is proportional to the cosine distance to its class centre and is also directly related to the face utility. In SER-FIQ, this method mitigates the need for any automated or human labelling by correlating the FIQ to the robustness of the extracted face embeddings to random dropout patterns. Face images with high utility are expected to have similar face representations resulting in low variance and vice versa.

**FaceQnet** [29] and **SDD-FIQA** [63] are categorised into the supervised FIQA methods. While FaceQnet uses ICAO compliant face images as high-quality reference images and builds quality labelling scores based on the FR genuine comparison score to ICAO compliant images, the SDD-FIQA uses pseudo quality label based on the Wasserstein distance between similarity distributions of genuine and imposter image pairs. Both training networks then fine-tuned a pre-trained FR base-network (RseNet-50 [64]) as the backbone architecture and the successive regression layer on top of the feature extraction layers are used to associate an input image to an FIQ score.

A ranking-based FIQA method learnt on comparing pairs of input images is **RankIQ** [30]. Chen et al. [30] are inspired by the idea that it is easier for a human operator to rank pairs in a relative manner than to define an absolute quality score for a single image. This method combines two-stage learning to first extract individual handcrafted features (e.g. HoG, Gabor, LBP, and CNN features) and later map these feature scores to a final quality score by using a kernel trick. Compared to the other considered FIQA methods, this method does not relate to very deep features or deep backbones, however, proved to perform competitively to recent FIQA approaches [65].

The presented methods in this section are designed for either assessing the general image quality or the face image utility (for FR) in the design concept. A number of the later methods were shown to predict low qualities for images where FR models focus on areas beyond the central face area [66], which might point out extreme poses among other factors. Our work investigates the face morphing effect on these measures. Table 1 listed the interpretation of these quality measures. The term ‘increasing’ indicates the ascending order of this quality score and means the higher the score the better the perceived quality or utility, and vice versa.

## 4 | EXPERIMENTAL DESIGN

The experiments in this work are designed so that they provide comprehensive answers to the following two research questions: (1) What is the effect of different morphing approaches on different IQA and FIQA measures, and (2) is it possible to use this effect to detect morphing attacks without the explicit training of an MAD? To address these two main investigations, we structured our experiment section first to introduce the variety of different morphing techniques and the 5 morphing datasets studied in this work. Each morph dataset intends to provide variations in the created morphing attacks depending on the bona fide source distributions extracted from the

**TABLE 1** The table shows the ordering of the quality measures considered in the manuscript. The term ‘increasing’ means that the way this quality measure is ordered is ascending and the higher the quality score, the better the perceived quality or utility and vice versa

FIQA					IQA				
MagFace [14]	SDD-FIQA [63]	SER-FIQ [62]	FaceQnet v2 [29]	RankIQ [30]	BRISQUE [17]	DBCNN [58]	CNNIQA [13]	UNIQUE [61]	dipIQ [36]
Increase	Increase	Increase	Increase	Increase	Decrease	Increase	Decrease	Increase	Increase

respective source face datasets. We subsequently introduce the evaluation metrics before presenting the protocols used in our investigations.

## 4.1 | Morphing techniques

The morphing techniques considered in this work can be grouped into classical landmark-based approaches and deep-learning-based approaches using generative adversarial networks (GANs). While the landmark-based morphing approaches could introduce visible blurring effects or other perceptual artefacts, the GAN-based approaches can generate more realistic morphing attacks. However, besides the difficult training procedure of GAN-based approaches, it is also more difficult to constrain the output of the GAN-based attacks. In the following, we will shortly explain these morphing techniques in their individual category.

### 4.1.1 | Facial landmarks-based morphing techniques

The **OpenCV** algorithm is an open-source implementation proposed by Satya et al. in Ref. [12] using the extracted locations of 68 facial landmarks determined by the Dlib library [67]. The facial landmarks are used to transform both bona fide source images into one morphed image using Delaunay triangles to warp the facial landmarks and alpha blending. This morphing approach has been proven to produce attacks leading to a relatively high vulnerability in FR systems [7, 39, 68].

**FaceMorpher** [69] is another open-source and facial landmark-based morphing algorithm, which is similar to the OpenCV algorithm. Instead of using Dlib, this method is based on STASM [70] as the facial landmark detector. These two landmark-based morphing algorithms create morphed face images with visible artefacts especially around the hair region because the regions outside the landmark positions are simply averaged.

**WebMorph** [71] is an online tool provided by the FRLI dataset provider to generate morphed images. This method leveraged 189 facial landmarks that are specifically labelled for the FRLI dataset. The high amount of precisely labelled landmark positions can generate morphs with reduced visible artefacts compared to using only 68 facial landmarks. However, as this method works exclusively for annotated FRLI images, this morphing technique is used only on the FRLI dataset.

**AMSL** is introduced by Neubert et al. in Ref. [9], which generates morphed images from bona fide images of FRLI dataset [72] using a private Combined Morphs tool. This tool enables generating morphed images with no visible ghosting artefacts around the hair and neck areas compared to OpenCV, FaceMorpher, and WebMorph, due to its post-processing step such as Poisson image editing.

### 4.1.2 | Generative adversarial network-based morphs

Leveraging the recent advances in deep-learning and generative adversarial networks (GANs), these GAN-based methods can generate more photo-realistic images without visible artefacts such as those blending artefacts introduced from other landmark-based morphing techniques. **StyleGAN 2** [73] is pre-trained on FFHQ dataset introduced in Ref. [74] and can generate high resolution and realistic faces without introducing noticeable artefacts. The StyleGAN was first used in Ref. [75] to produce morphed images, which later used an updated generator in Ref. [7].

**MorGAN** [49] was the first proposed generative approach. MorGAN uses a specific loss function so that instead of generating a morph with most probable lost identity using the GAN training, the network enforces the generated morphs to keep the identity information. However, the output of the MorGAN generated morphs are relatively unrealistic with a resolution of  $64 \times 64$  pixels.

## 4.2 | Datasets

Five different and diverse morphing datasets are investigated in this work to provide generalised conclusions. They are composed of FRLI-Morphs [68, 72], FERET-Morphs [68, 76], and FRGC-Morphs [68, 77], LMA-DRD [39], and MorGAN [49].

Research in the field of face MAD becomes more stringent, but there is a lack of shared, publicly available datasets. To overcome this gap, Sarkar et al. [68] worked on creating publicly available morphing datasets with several forms of attacks to be used for research purposes. They provided new morphing datasets with five different types of morphing attacks based on OpenCV [12], FaceMorpher [69], WebMorph [71], AMSL [9] and StyleGAN 2 [73]. The attacks are created using source face images from three publicly available face datasets including the Face Research London Lab (FRLI) dataset, the Facial Recognition Technology (FERET) dataset, and the Face Recognition Grand Challenge (FRGC) dataset. These three morph datasets also build a big part of the foundation of our research in this work.

The **FRLI-Morphs** dataset [68] is generated from the publicly available Face Research London Lab dataset [72]. This dataset includes 5 different morphing techniques, including OpenCV [12], FaceMorpher [69], StyleGAN 2 [73], WebMorph [71] and AMSL [9]. Each morph technique contains 1222 morphed face images generated using only frontal face images with high resolution. We note, however, that the AMSL generates two unique morphed images for every pair of source bona fide images, which doubles the size of the AMSL morphs compared to other morphing techniques. The bona fide face images are extracted using the provided protocol from the FRLI-Morphs dataset [72] using only the frontal and smiling face images. Even though the source images are of very high

visual quality and under uniform illumination with a large variety of ethnicity, pose, and expression, the number of bona fide samples in the dataset is limited to 204 images from 102 identities.

The **FERET-Morphs** dataset [68] and the **FRGC-Morphs** dataset [68] are extracted from the official FERET [76] and FRGC v2.0 dataset [77], respectively. Both datasets include 3 different morphing techniques, including OpenCV, FaceMorpher, and StyleGAN 2. For each morphing technique, it contains 529 and 964 morphed face images for FERET-Morphs and FRGC-Morphs, respectively. Opposite to FRL, the FERET and FRGC v2.0 dataset contained face images of large variability in quality and thus enabling creating morphing attacks with larger quality distributions compared to FRL-Morphs as can be seen in Figure 3. The bona fide images are extracted using the provided protocols in the FERET-Morphs and FRGC-Morphs datasets.

The **LMA-DRD** Morph dataset [39] contains morphed face images created from the VGGFace2 [78] dataset. The images used to generate morphs are frontal, with a neutral expression, and according to the generation details described in Ref. [39]. The large VGGFace2 allows the selection of high-quality images according to ICAO standards. The morphing technique is based on OpenCV using facial landmarks. For further details of the morphing process and the parameters used, we refer to the work in Ref. [39]. An additional bona fide image is selected for each morphed identity, when available. In total, the used LMA-DRD morphing dataset contains 276 digital Bona fide (D-BF) images and 364 digital morphing attacks (D-M). These images were printed on  $11.5 \times 9$  cm glossy photo paper in a professional studio and scanned with a 600 dpi scanner. They resulted in the same number of re-digitised bona fide (PS-BF) and attacks (PS-M), leading to two versions of the dataset, the digital LMA-DRD (D) and the re-digitised LMA-DRD (PS). Opposite to other morphing datasets, the LMA-DRD dataset specifically aims at investigating the morphing attacks after the re-digitisation process. The re-digitised versions are intended to simulate real identity document issuing scenarios and their effect on the digital artefacts, which simulate a realistic challenge for MAD algorithms.

To enable selection of similar faces to generate more sophisticated morphs, the **MorGAN** dataset [49] created morphing attacks from a large pool of identities using CelebA [79]. The CelebA dataset contains around 202,599 face images of 10,177 identities, which allows the creation of morphing attacks fulfilling the intended goal. The MorGAN dataset contains 1000 morphed images for each of the two morphing attacks. One attack form is based on MorGAN and the other is created using the landmark-based morphing approach using OpenCV [12]. The bona fide images contain two sets (1500 bona fide references and 1500 bona fide probes). The exact description of the morphing dataset can be found in Ref. [49].

In Table 2, the presented 5 morphing datasets with their respective source datasets and the included morphing techniques are listed. In addition, the number of the used images in the individual setting is provided along with the bona fide image numbers and the total amount of images.

For all images across the morphing datasets, the images are pre-processed using the MTCNN framework [80] to detect, crop, and align (geometric transform) the face to output size of  $224 \times 224$  pixels whenever necessary. The face images in the MorGAN dataset are already pre-processed with an output size of  $64 \times 64$  pixels. For each of the attack forms across the 5 morphing datasets (in total 15) a pair of example images are depicted in Figure 1. Figure 2 further visualises bona fide face image pairs for the 5 morphing datasets, where the LMA-DRD dataset has two different bona fides, one for the source digital images and the other for the re-digitised (PS) images.

Table 3 further list the quality scores for the example images selected in both figures with 3 best performing quality measures from both the FIQA and IQA categories. For all quality measures, a higher quality score indicates a better quality of the face image, except for BRISQUE and CNNIQA, which have the inverted meaning. For most quality measures listed in the table, it is observed that the bona fide images have a higher quality compared to the morphing attacks across most morphing datasets.

### 4.3 | Performance metrics

As targeted to answer both (1) the effect of morphing on quality and (2) the detectability of morphing attacks by probing this effect, we grouped the used evaluation metrics into these two individual classes.

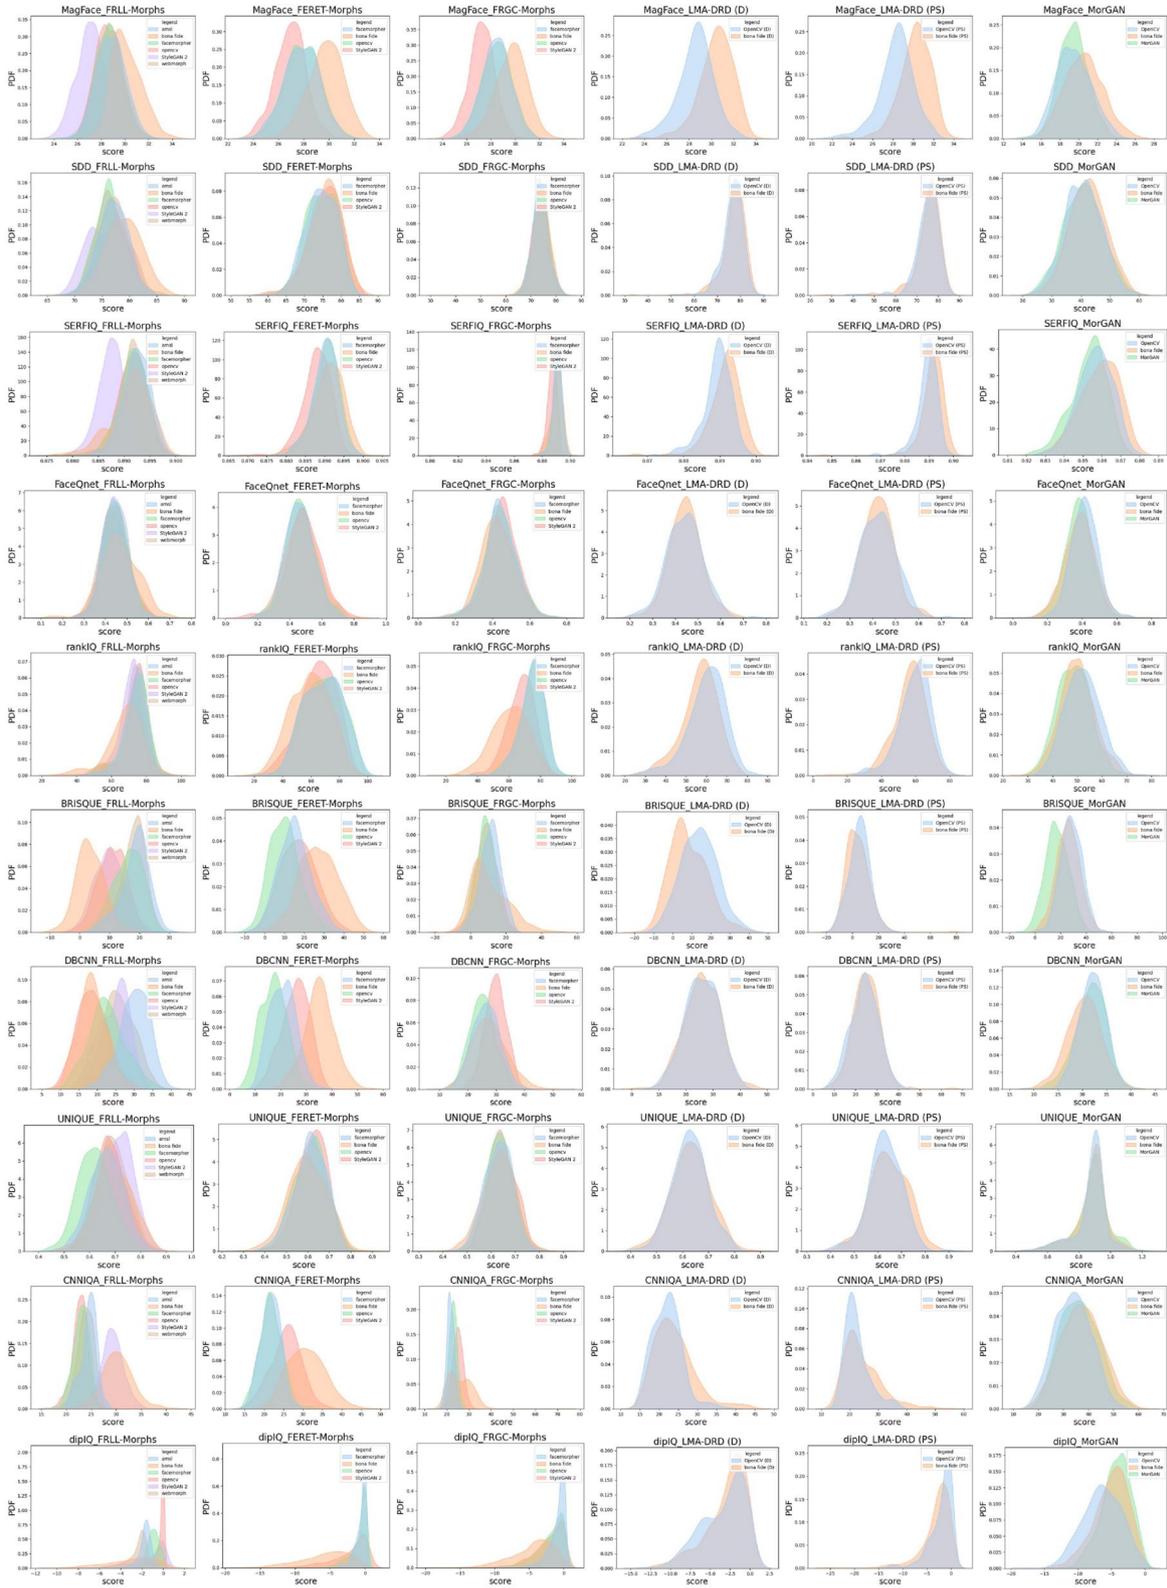
First, to consider the morphing attack and its effect on quality scores, we listed the quality scores for sample bona fide and attack images for reference. The different morphing dataset has different quality distribution. We show this by visualising the quality score distribution approximated by the kernel density estimation method on the discrete quality scores of (1) the bona fide images across morphing datasets or (2) for bona fide images and the morphing attacks within the same morphing dataset.

Another property to review is the separability between the quality distributions of morphing attacks and bona fide. The Fisher Discriminant Ratio (FDR) is a useful measure to do just that. We looked at the quality score distributions of the individual attacks and the bona fide images within the same morphing dataset. The term FDR in Ref. [81] is defined by Equation (1) and quantifies the separability of the score distributions of morphed and bona fide face images.

$$FDR = \frac{(\mu_1 - \mu_2)^2}{(\sigma_1)^2 + (\sigma_2)^2}, \quad (1)$$

In Equation (1), the terms  $\mu_1$  and  $\mu_2$  are the mean quality scores of the morphed and bona fide distributions, respectively, and  $\sigma_1$  and  $\sigma_2$  describe the standard deviation term of both distributions, respectively.

To draw conclusions on the detectability of the morphing attack by quality measures, we use the presentation attack detection measures defined in ISO/IEC 30107-3 [82] as the



**FIGURE 3** Figure visualises the distributions of the different morphs and the bona fide scores based on the investigated quality measures for the 5 morphing datasets, where the LMA-DRD further splits into the digital and re-digitised versions

performance metrics in our work. We used the terms (1) Attack Presentation Classification Error Rate (APCER) describing the proportion of attack presentations wrongly

classified as bona fides under certain morphing scenarios, (2) the Bona Fide Presentation Classification Error Rate (BPCER) describing the proportion of the error of bona fide face

Morph dataset	Source dataset	No of images			Total
		No. Bona fide	Morphing type	No. attack	
FRLM-morphs [68]	FRLM [72]	204	AMSL [9]	2444	2648
			FaceMorpher [69]	1222	1426
			OpenCV [12]	1222	1426
			StyleGAN 2 [73]	1222	1426
			WebMorph [71]	1222	1426
FERET-morphs [68]	FERET [76]	1413	FaceMorpher [69]	529	1942
			OpenCV [12]	529	1942
			StyleGAN 2 [73]	529	1942
FRGC-morphs [68]	FRGC v2 [77]	3167	FaceMorpher [69]	964	4131
			OpenCV [12]	964	4131
			StyleGAN 2 [73]	964	4131
LMA-DRD (D) [39]	VGGFace2 [78]	276	OpenCV (D) [12]	364	640
LMA-DRD (PS) [39]	VGGFace2 [78]	276	OpenCV (PS) [12]	364	640
MorGAN [49]	CelebA [79]	3000	OpenCV [12]	1000	4000
			MorGAN [49]	1000	4000

**TABLE 2** The used morphing datasets are depicted in the table including their source dataset and the morphing techniques used in the individual dataset to create the morphing attacks. In addition, the number of the used images under each setting is also provided

images, which are wrongly classified as morphed face images, and (3) the Equal Error Rate (EER) representing the *APCER* or *BPCER* at the operation point where the *APCER* and *BPCER* are the same value. We further provide the Average Classification Error Rate *ACER* describing the average error between *APCER* and *BPCER* as given in Equation (2).

$$ACER = \frac{APCER + BPCER}{2} \quad (2)$$

The ACER metric is not part of Ref. [82] standard anymore but helps in giving a single value indicative of the performance. A lower error rate, for all used error metrics, indicates a higher PAD performance.

#### 4.4 | Investigations

The experiments are designed to address the two main research issues: (1) the effect of different morphing approaches on quality and (2) the possibility of using these effects to detect morphing attacks.

For the first aspect, the effect of morphing on quality, we (1) investigate the bona fide of each dataset versus each of the morphing techniques of this dataset. This investigation includes a visual presentation of the quality (different measures) distributions of the bona fide of each dataset alongside the morphing approaches in these datasets. The investigation also includes (2) presenting a quantitative measure of how different are morphing attacks from the bona fide images in each dataset, that is, the separability FDR measure. For the second aspect, the possibility of using the morphing effect on quality

to detect attacks, we (1) first look at the possibility to maintain the correct detection of bona fide, even if the bona fide comes from a different source. To do that, we fix the BPCER on each data source and calculate the BPCER on all other data sources. Then, (2) we investigate to use this BPCER threshold on each data source to detect the attacks (APCER), despite their differences, in that dataset and all other datasets. Third (3), after having a detailed look at detecting bona fide and attacks separately (BPCER and APCER), we derive a final overall conclusion by reporting the detailed ACER measures on all datasets and attacks, while fixing the BPCER at each data source individually.

## 5 | RESULTS AND DISCUSSION

The two main research gaps that this work addresses are (1) the effect of face morphing on image quality and utility and (2) the detectability of these different morphing attacks based on image quality and utility. The experimental outcome regarding these two issues is discussed based on the detailed results in this section.

### 5.1 | The effect of morphing on quality

#### 5.1.1 | How different are the quality score distributions of bona fide and morphed samples?

To answer this question, we look at the quality score distributions produced by each of the 10 used quality measures. The distribution comparison is within each dataset bona fide and

**TABLE 3** Quality scores of the displayed sample images from Figures 3 and 4 are listed in here. A higher score indicates better quality except for BRISQUE and CNNIQA, which possess a reversed ordering. This results in the observation of bona fide images having better quality compared to most of the morphing attacks across most morphing datasets, especially in the case of FRLL-Morphs, FERET-Morphs, and FRGC-Morphs

Dataset	Morph-type	FIQA						IQA					
		MagFace		SER-FIQ		RankIQ		BRISQUE		CNNIQA		dipIQ	
		Left	Right	Left	Right	Left	Right	Left	Right	Left	Right	Left	Right
FRLL-Morphs	AMSL	26.9800	26.9126	0.8943	0.8928	68.6190	73.8720	12.9910	14.0530	23.4542	22.6041	-3.1760	-3.3660
	FaceMorpher	27.1052	26.3664	0.8919	0.8889	71.9610	61.8510	10.3640	10.6400	19.1544	17.9703	-0.4410	-1.1200
	OpenCV	27.3879	27.8763	0.8920	0.8928	74.0580	57.1580	8.8580	7.1440	18.4415	17.0790	0.0000	-0.1540
FERET-Morphs	StyleGAN 2	26.9321	27.5039	0.8916	0.8878	69.9110	70.2580	6.0070	8.8480	21.5867	19.5706	-0.5050	-1.8630
	WebMorph	27.0006	27.5805	0.8912	0.8910	74.1640	59.0160	17.8020	14.3810	22.4474	19.8183	-2.7940	-4.4660
	Bona fide	28.4257	30.0034	0.8933	0.8930	74.5040	56.9300	-0.1820	0.7070	20.3228	20.6871	-2.7890	-2.6990
FRGC-Morphs	FaceMorpher	27.6738	26.1388	0.8907	0.8870	54.4640	65.0830	12.3540	9.0630	19.3835	18.1728	-0.6580	-0.5450
	OpenCV	26.8965	27.1316	0.8878	0.8870	62.7130	67.9110	7.3530	6.1690	19.0922	17.9139	0.0000	0.0000
	StyleGAN 2	25.9665	26.8775	0.8859	0.8876	59.3050	48.5360	25.0580	15.4620	24.3433	21.2487	-5.1980	-0.0590
LMA-DRD (D)	Bona fide	29.4893	29.5968	0.8937	0.8952	60.2820	62.6590	43.9160	33.9450	39.5276	34.0680	-1.3870	-1.3670
	FaceMorpher	28.0497	27.0605	0.8904	0.8859	71.7660	60.9050	6.0720	6.5410	15.0684	19.3062	-2.4470	-0.4390
	OpenCV	27.4370	27.8318	0.8903	0.8854	69.5100	64.9240	4.3500	3.5600	17.1487	15.6377	-3.0620	-2.0350
LMA-DRD (PS)	StyleGAN 2	27.6772	29.1389	0.8841	0.8919	62.4120	56.3640	2.7130	10.11530	18.7467	17.8662	-2.8610	-3.6840
	Bona fide	28.6661	28.8975	0.8925	0.8879	54.6710	49.4110	5.7000	10.9440	17.4571	19.5861	-9.6210	-1.4540
	OpenCV (D)	29.0780	27.0214	0.8901	0.8918	70.4170	61.3760	7.0480	18.8090	15.8341	24.1678	-0.7470	-5.3790
MorGAN-morphs	Bona fide (D)	28.7873	29.9078	0.8913	0.8905	64.4650	51.9640	0.8200	-2.5890	17.5215	21.2466	0.0000	-0.1550
	OpenCV (PS)	29.1056	27.0557	0.8906	0.8924	73.2130	67.9370	1.2400	5.9050	14.6985	15.8074	-2.7770	-0.0440
	Bona fide (PS)	28.9074	29.9588	0.8928	0.8913	69.5030	52.0630	-1.6830	-6.2090	17.5302	16.4102	-1.9700	-2.0260
MorGAN	OpenCV	19.7456	14.7414	0.8557	0.8478	44.3260	44.6990	25.7080	18.0130	20.9167	22.0238	-5.5650	-7.8220
	Bona fide	17.6934	18.1678	0.8510	0.8515	47.4040	48.4310	10.6440	-0.0210	29.4583	28.5578	-5.5000	-5.5400
	OpenCV	16.7996	23.6150	0.8472	0.8630	35.7060	51.4440	-1.2140	36.9810	30.8734	26.1947	-3.0170	-2.6580

morphing attacks. Figure 3 shows the quality score distributions of different morphing approaches versus its bona fide distribution within each of the 5 morphing datasets. Each row represents one IQ or FIQ measure. Observing the quality of morphing attacks and bona fide as score distributions in Figure 3, we see that different morphing attacks shift the value of IQ/FIQ compared to bona fide images for most quality measures.

In the FRLL-Morphs, landmark-based attacks showed closer quality distribution to the respective bona fide images, while StyleGAN 2 generated morphs clearly showed lower quality compared to landmark-based morphs measured by quality measures such as MagFace, SDD-FIQA, SER-FIQ, and BRISQUE. The IQA measures such as UNIQUE, CNNIQA, and dipIQ even signed slightly higher quality for StyleGAN 2 generated morphing attacks than bona fide distributions. FaceQnet seems to have more difficulty distinguishing the quality difference between the morphs and bona fide images. Similar conclusions are to be drawn for both FERET-Morphs and FRGC-Morphs, where StyleGAN 2 created morphs show lower image quality measured by MagFace, SERFIQ, BRISQUE, and DBCNN. In FERET-Morphs, DBCNN assigns the bona fide images better quality compared to other morphing attacks of different kinds. CNNIQA on the contrary assigns lower quality values to bona fide images compared to morphing attacks.

For digital and re-digitised attacks in the LMA-DRD dataset, only MagFace, SER-FIQ, RankIQ, and BRISQUE clearly show a quality degradation of morphing attacks towards bona fide images. CNNIQA even seems to have slightly higher quality for the digital and re-digitised morphing attacks

compared to bona fide. Looking at the MorGAN dataset, consistent findings as in Ref. [10, 49] are observed, where we could confirm that IQA measures, such as CNNIQA and dipIQ have similar image quality between MorGAN generated attacks and bona fide, while BRISQUE shows the opposite conclusion. SER-FIQ and BRISQUE put the quality of the landmark-based attacks closer to the bona fide images. However, as it is noticed in Figure 3, the quality of MorGAN attacks and bona fide, for almost all FIQA and IQA, is lower than other datasets (keep in mind that some quality measures has reversed interpretation as presented in Table 1). As bona fide samples of the MorGAN data are also of relative low quality, this must be strongly influenced by the low resolution of MorGAN data ( $64 \times 64$  pixels). This strong influence of resolution is evident as it is the only major factor that differentiates between the bona fide samples of all other datasets and the ones of the MorGAN dataset.

### 5.1.2 | Are the quality distributions of bona fide and attack samples separable? And how separable are they?

Table 4 shows the FDR value between each morphing attack to the bona fide quality distribution within each morphing dataset. A large FDR value indicates a high separability between both distributions and thus indicates the possibility of using this quality metric for detecting this kind of morphing attack among the bona fide images.

One finding based on the FDR in Table 4 is that we observe MagFace having relatively high separability for GAN-

**TABLE 4** FDR scores indicate the separability between morphs and bona fide distributions based on the quality metric. Only intra-dataset comparison is considered here. A larger FDR score indicates a better separability between the morphing attack and bona fide distributions

Dataset	Morph-type	FIQA					IQA				
		MagFace	SDD-FIQA	SER-FIQ	FaceQnet	RankIQ	BRISQUE	DBCNN	CNNIQA	UNIQUE	dipIQ
FRLL-morphs	AMSL	0.5942	0.2021	0.0546	0.0506	0.2063	<b>5.1530</b>	4.1005	2.7154	0.0737	0.2901
	FaceMorpher	0.7354	0.3761	0.0052	0.0294	0.2157	3.5891	0.3924	<b>4.3623</b>	0.5835	0.8133
	OpenCV	0.8739	0.2959	0.0079	0.0450	0.2092	1.2141	0.0005	<b>4.6174</b>	0.0270	1.4129
	StyleGAN 2	<b>3.6564</b>	0.5017	0.6746	0.0246	0.0462	0.9034	1.0804	0.0680	0.0744	0.3583
	WebMorph	0.8103	0.1516	0.0014	0.0466	0.1146	<b>5.7858</b>	1.6706	3.3614	0.0014	0.0532
FERET-morphs	FaceMorpher	0.9094	0.0281	0.0148	0.0092	0.2002	0.9251	<b>3.3444</b>	2.8108	0.0098	2.1001
	OpenCV	0.8483	0.0197	0.0058	0.0057	0.2071	1.9369	<b>5.2332</b>	2.6064	0.0138	1.7226
	StyleGAN 2	<b>2.2707</b>	0.0000	0.2686	0.0082	0.0760	0.5868	1.3572	0.6977	0.0448	0.6578
FRGC-morphs	FaceMorpher	0.4877	0.0021	0.0064	0.0122	0.9074	0.0009	0.0280	<b>0.6465</b>	0.0001	1.2222
	OpenCV	0.4492	0.0015	0.0102	0.0109	<b>0.8586</b>	0.0428	0.1082	0.3502	0.0009	0.5369
	StyleGAN 2	<b>2.1508</b>	0.0001	0.2715	0.0278	0.2116	0.0242	0.0855	0.0388	0.0239	0.3793
LMA-DRD (D)	OpenCV (D)	<b>0.7061</b>	0.0019	0.1309	0.0010	0.0654	0.1646	0.0002	0.0767	0.0017	0.0112
LMA-DRD (PS)	OpenCV (PS)	<b>0.6552</b>	0.0006	0.1144	0.0001	0.0423	0.0102	0.0075	0.0648	0.0201	0.0668
MorGAN	OpenCV	<b>0.1562</b>	0.0291	0.0279	0.0375	0.0466	0.0350	0.1349	0.0937	0.0005	0.1157
	MorGAN	<b>0.1855</b>	0.0268	0.1619	0.0055	0.0194	0.4380	0.0728	0.0050	0.0015	0.0442

based morphing attacks (such as StyleGAN 2, MorGAN, and even for re-digitised morphing attacks in the LMA-DRD dataset), while RankIQ, BRISQUE, and CNNIQA outperform mostly in separating landmark-based morphing attacks. The highest FDR values in Table 4 are marked in bold for improved visibility. This finding might be justified by the nature of the IQA measures looking at the general image quality but having trouble focussing on learning the specific GAN-based artefacts in a face image. BRISQUE and CNNIQA as representative quality measures of the IQA family focus more on the visible artefacts in the morphed images and less on realistic artefacts introduced by GAN-based approaches. These visible artefacts are easier to be introduced by landmark-based morphing techniques using alpha blending of two bona fide sources such as in the case of FaceMorpher, OpenCV, and WebMorph in the FRL-Morphs, FERET-Morphs, and FRGC-Morphs datasets as visualised in Figure 1. Another CNN-based quality measure using multitask learning (DBCNN) even outperforms CNNIQA in FERET-Morphs, showing stronger separability for landmark-based morphs to bona fide as well as for StyleGAN 2 generated morphs (but still inferior to MagFace in this case).

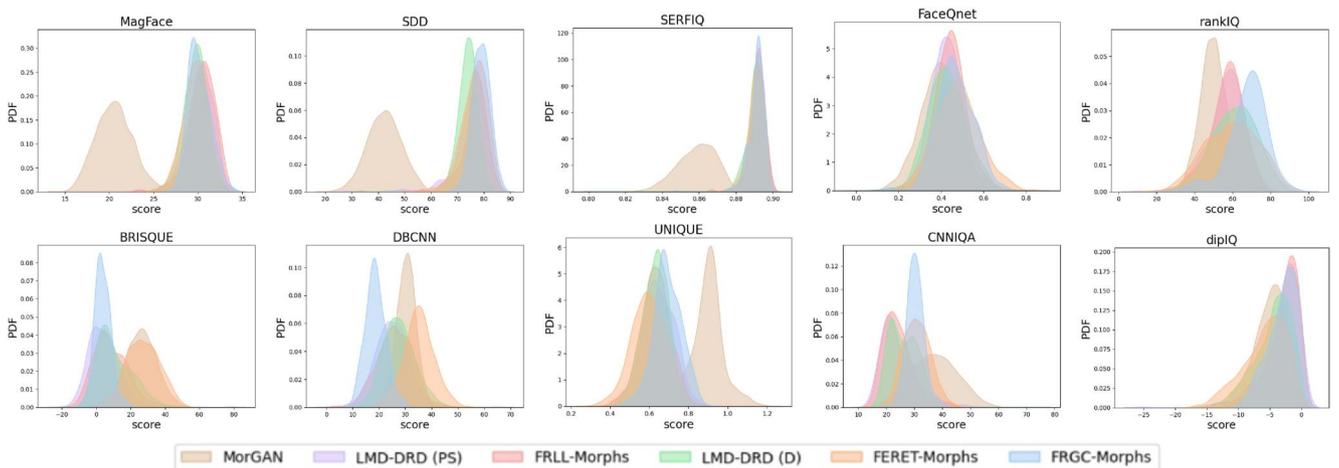
### 5.1.3 | Are morphing attacks with relatively low/high IQ also lead to low/high FIQ?

Several studies in Ref. [34, 65] pointed out that the FIQ measures such as MagFace [14], SDD-FIQA [63], SER-FIQ [62], and CR-FIQA [55] are highly correlated with the face utility as defined in ISO/IEC 29794-1 [83]. A recent work [65] has shown that for normal FR samples, IQ measures also correlate to utility but to a much lower degree than FIQ. That work also showed low correlation between the decisions of IQA and FIQA. Knowing that IQA intends to measure the perceptual quality and FIQA intends to measure the image utility, one can see from the quality distributions in Figure 3

that not only the IQ is affected by morphing but also the utility related FIQ measures, more clearly for MagFace than others.

To quantitatively represent the correlation between IQA and FIQA, we look into the lowest and highest quality samples for a set of FIQA and IQA measures and measure the ratio of the shared (overlapped) samples in each of these groups (high or low quality). To calculate this correlation between quality estimation methods, we calculate the attack samples overlap ratio between the samples of the lowest quality (10% of the data) between every pair of quality estimation methods and the same for the 10% of the highest quality. A large overlapping ratio indicates a larger reasoning similarity between the considered pair of methods. These ratios are presented in Figure 5 for the top 10% qualities and in Figure 6 for the bottom 10% quality, for the 2 FIQA and 2 IQA methods that showed the largest separability between morphing attacks and bona fide samples, that is, MagFace and RankIQ from FIQA methods, and CNNIQA and dipIQ from IQA methods. For most morphing attacks, the correlation is slightly larger within FIQA methods (i.e. between MagFace and RankIQ) and within IQA methods (i.e. CNNIQA and dipIQ) than it is between FIQA and IQA methods. This is not always true, as CNNIQA and MagFace show a relatively higher correlation when assessing StyleGAN2 morphs. However, the correlation between IQ and FIQ measures is relatively low, indicating that, as proven previously on bona fide samples in Ref. [65], an image with low perceptual quality does not necessarily lead to a low FR utility and vice versa. A clear example of that is applying FR on masked faces [84, 85], where an image of a masked face can be of high perceptual quality, but it is of low utility given the mask occlusion [57].

So far, our investigations showed that some quality measures are clearly affected by the morphing process. We have seen also that the different morphing techniques affect these quality measures differently. A major outcome is that some quality measures, such as MagFace, do result in a high separability between bona fide samples and the different attacks. This will be leveraged in the next section to investigate how



**FIGURE 4** The figure depicts the bona fide quality score distributions using kernel density estimation approximation of the 5 morphing datasets according to the proposed quality measures. MorGAN bona fide scores reveal lower quality compared to other bona fide images from other morphing datasets. This figure also provides a slight hint of the quality of the underlying source datasets from which the morphing attacks are created

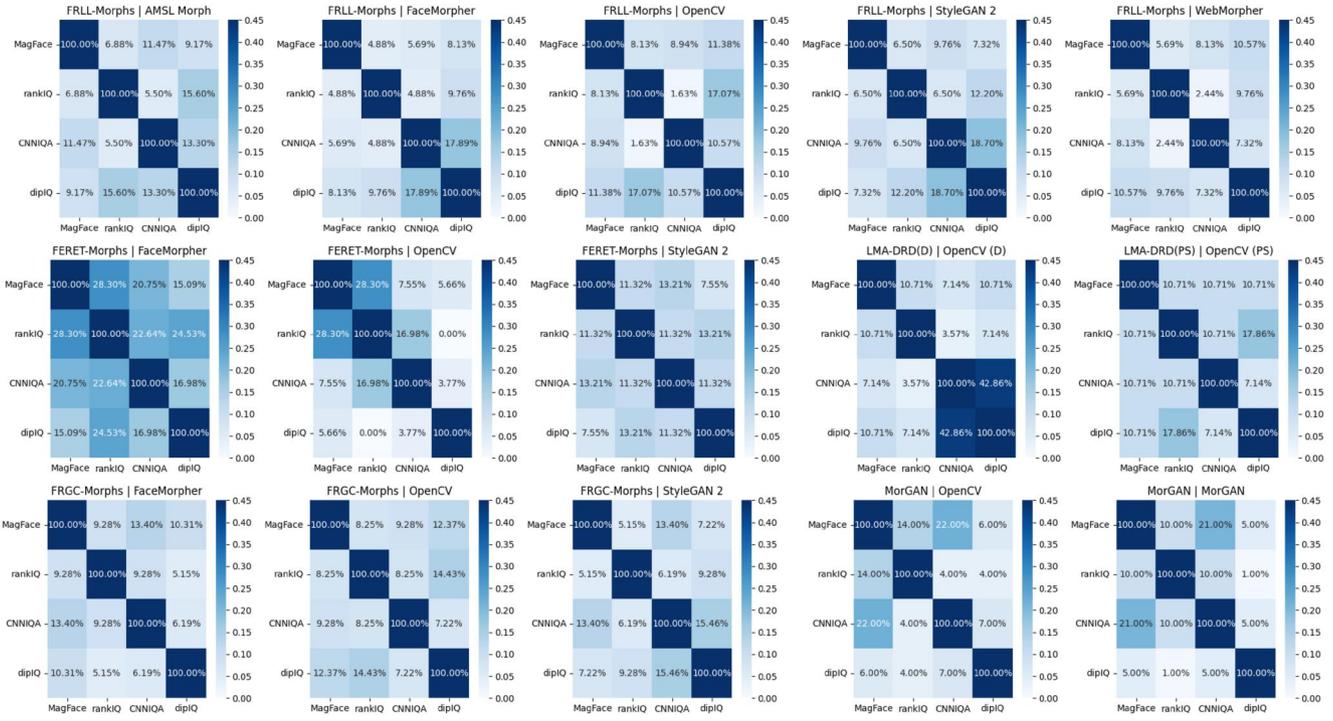


FIGURE 5 The overlap ratio between the morphing attack samples designated to be among the highest 10% qualities (in their respective datasets) as ranked by different pairs of FIQA and IQA methods (on the  $X$  and  $Y$  axes). MagFace and RankIQ are FIQA methods, and CNNIQA and DeepIQ are IQA methods. This indicates the morphing attacks with relatively high IQ do not necessary lead to high FIQ

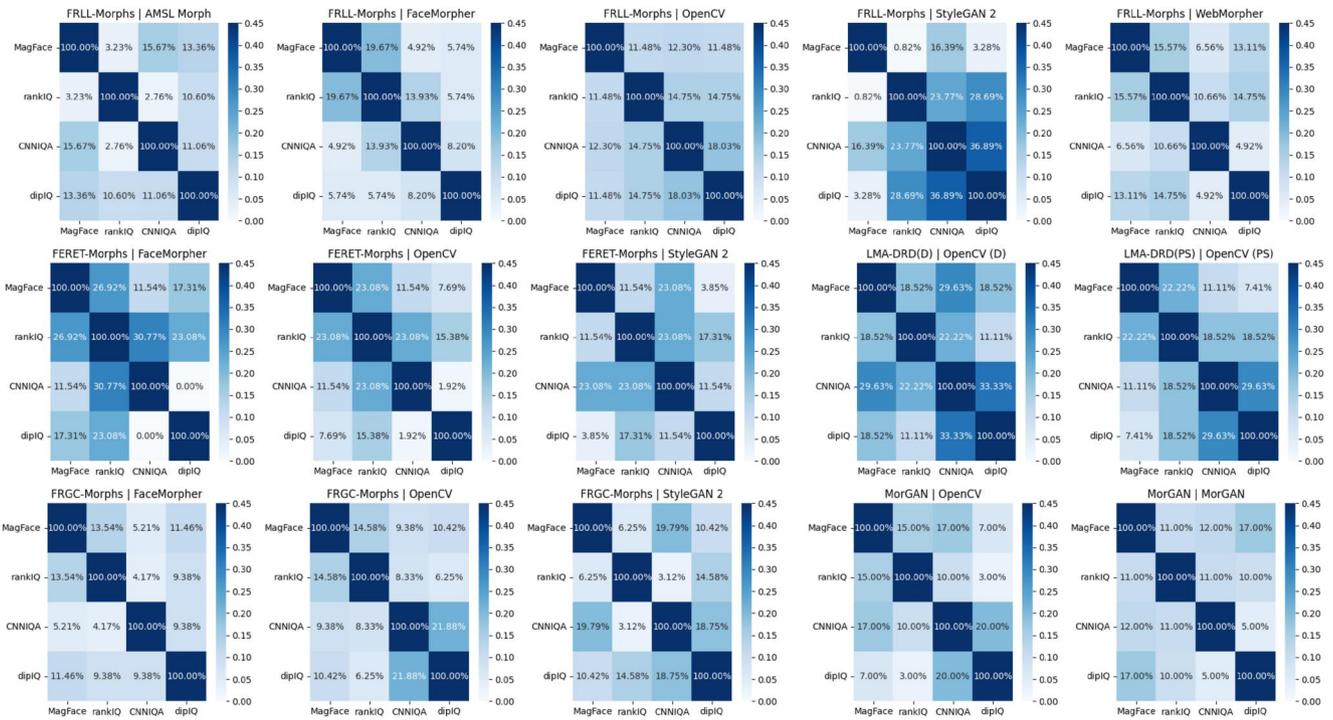


FIGURE 6 The overlap ratio between the morphing attack samples designated to be among the lowest 10% qualities (in their respective datasets) as ranked by different pairs of FIQA and IQA methods (on the  $X$  and  $Y$  axes). MagFace and RankIQ are FIQA methods, and CNNIQA and DeepIQ are IQA methods. This indicates the morphing attacks with relatively low IQ do not necessary lead to low FIQ

can this separability be used to detect morphing attacks and how generalisable is this detection over multiple morphing attack methods and bona fide sources. A side observation is that morphing attacks with relatively low/high IQ does not necessarily lead to low/high FIQ, confirming previous studies' conclusions on normal FR tasks [65].

## 5.2 | The unsupervised detectability of morphing attacks by quality measures

In this section, we start by looking at the detectability within each dataset, motivated by the bona fide-attack separability presented in the last section. Then, motivated by the differences in the quality distributions of the different bona fide sources, we investigate the generalisability of detecting morphing attacks/bona fide by quality over different bona fide sources and attack methods.

### 5.2.1 | If quality measures were used as an MAD detection score, what would be the detection error within each of the investigated datasets?

Table 5 shows the EER for each morphing technique in each dataset by using the quality measures as MAD scores. It is to

note  $\star$  behind the quality and utility measure in the table. This  $\star$  indicates the reversed ordering of the quality measure to its original meaning as listed in Table 1. Looking at the quality distribution for dipIQ, we observed that the bona fide quality is lower compared to the morphing attacks. Therefore, in order to use this quality metric as an MAD score, we have to reverse the quality scores to make the score thresholds consistent with other quality measures considered in this work. The same is for RankIQ, DBCNN, and CNNIQA measures. The smaller the EER, the better is this quality metric suited for MAD task. The numbers in bold indicate the lowest EER for each attack type across all used quality measures within the same morphing dataset. The last row in the table shows the *mean* overall EER scores for each quality metric across all morphing types and datasets. This overall performance provided us the indication, which quality metric performs generally well on most of the MAD tasks. The best 3 performing quality measures from the FIQA and IQA categories are selected for further detailed cross-dataset investigations. Therefore, for the rest of the experiments, 6 FIQA/IQA methods will be considered, namely the MagFace, SER-FIQ, RankIQ, BRISQUE, CNNIQA, and dipIQ.

Based on the EER values in Table 5 and consistent with the previously discussed FDR values, MagFace shows the lowest EER for StyleGAN 2 generated morphs in the FRLMorphs, while inverted CNNIQA shows the lowest EER for

**TABLE 5** EER score is presented between the morphs and bona fide distributions using the quality metric as a decision threshold. The quality measures with \* indicate that using the inverse of the quality value (lower quality indicates bona fide) leads to a lower average EER and thus is used. Each EER operation point is calculated using the morphing attack and the bona fide scores of the same dataset. In each row of the table, the quality metric with the lowest EER between bona fide and morphing attacks within the same morphing dataset is marked in bold. Only the *mean* EER value in the last row is averaged across the morphing attacks and morphing datasets alongside the quality measure to determine the best performing 3 measures from both the IQA and FIQA categories, which are in bold

Dataset	Morph-type	FIQA					IQA				
		MagFace	SDD-FIQA	SER-FIQ	FaceQnet	RankIQ*	BRISQUE	DBCNN*	CNNIQA*	UNIQUE	dipIQ*
FRLMorphs	AMSL	0.3122	0.3908	0.5471	0.4524	0.3646	0.0694	0.0883	<b>0.0791</b>	0.4437	0.3572
	FaceMorpher	0.2856	0.3232	0.4935	0.4755	0.3519	0.1023	0.3077	<b>0.0360</b>	0.2807	0.2610
	OpenCV	0.2645	0.3620	0.4881	0.4644	0.3407	0.2187	0.5029	<b>0.0541</b>	0.4758	0.1384
	StyleGAN 2	<b>0.0704</b>	0.3437	0.2365	0.4640	0.4272	0.2570	0.2381	0.4313	0.5859	0.3101
	WebMorph	0.2826	0.3948	0.4799	0.4480	0.3956	0.0541	0.1867	<b>0.0713</b>	0.5061	0.4308
FERET-Morphs	FaceMorpher	0.2420	0.4386	0.4575	0.4688	0.3913	0.7391	0.8941	<b>0.1080</b>	0.5425	0.1040
	OpenCV	0.2514	0.4480	0.4688	0.4669	0.3913	0.8431	0.9452	<b>0.1229</b>	0.5501	0.1512
	StyleGAN 2	<b>0.1399</b>	0.4915	0.3459	0.5198	0.4178	0.6957	0.7977	0.2892	0.5728	0.2647
FRGC-Morphs	FaceMorpher	0.3039	0.4761	0.5062	0.5353	<b>0.2417</b>	0.4305	0.5280	0.3094	0.4907	0.1815
	OpenCV	0.3143	0.4793	0.5156	0.5353	<b>0.2448</b>	0.5031	0.5820	0.3703	0.4948	0.2977
	StyleGAN 2	<b>0.1432</b>	0.4948	0.3268	0.5685	0.3662	0.4886	0.4004	0.4741	0.5207	0.3195
LMA-DRD (D)	OpenCV (D)	<b>0.2428</b>	0.4783	0.3659	0.5000	0.4239	0.3913	0.4964	0.4529	0.4783	0.5362
LMA-DRD (PS)	OpenCV (PS)	<b>0.2509</b>	0.4727	0.3709	0.5091	0.4473	0.4436	0.5127	0.4509	0.4691	0.3891
MorGAN	OpenCV	<b>0.3880</b>	0.4540	0.4470	0.5530	0.4400	0.4520	0.3940	0.4200	0.4970	0.6100
	MorGAN	<b>0.3750</b>	0.4560	0.3900	0.4780	0.5300	0.6820	0.4150	0.4770	0.4980	0.4540
<b>Mean</b>		<b>0.2577</b>	0.4335	<b>0.4293</b>	0.4959	<b>0.3849</b>	<b>0.4247</b>	0.4859	<b>0.2778</b>	0.4937	<b>0.3203</b>

all landmark-based morphs within this dataset. The corresponding finding is to be observed for FERET-Morphs. Regarding the FRGC-Morphs, the inverted RankIQ surprisingly shows lower EER for FaceMorpher and OpenCV morphs compared to the inverted CNNIQA. As for the LMA-DRD datasets and the MorGAN dataset, MagFace constantly outperforms the other quality measures in terms of EER by a large margin to further consolidate its superior ability to be used as a decision metric in the detectability of morphing attacks. We observed that MagFace and the inverted CNNIQA show overall low EER across all morphing attacks and datasets.

### 5.2.2 | How different are the qualities of different bona fide samples? And how does that effect the correct detectability of bona fide?

Figure 4 depicts the quality distributions of all the bona fide images across the 5 data sources using all 10 selected quality measures. Since the MorGAN dataset has been originated from lower image resolution, the quality of this dataset is visibly lower compared to other bona fide sources. This is shown by all FIQA measures. MagFace, SDD-FIQA, and SER-FIQ clearly placed the quality of the MorGAN bona fide images that are considerably lower than bona fide images in other data sources. Only CNNIQA, DBCNN, and UNIQUE did not place the MorGAN bona fide images as far from other bona fide images. Besides bona fide images in the MorGAN dataset, other datasets showed slightly different quality distributions, however, not consistently over all quality measures. FRLL bona fide images consistently showed tight quality distributions in comparison to other data sources, which might be due to the highly controlled capture environment in FRLL. The FERET bona fide samples showed relatively shifted quality distributions in comparison to other data sources, especially with the BRISQUE, DBCNN, and CNNIQA measures, which might be due to the fact that the FERET dataset was captured in the early 1990s with relatively outdated cameras and thus differences in the resulting images. In general, putting MorGAN (low resolution and thus low utility) aside, all bona fide sources showed similar FIQA distributions despite the differences in the perceived quality.

Given that the quality values can be used as an MAD score, setting the BPCER value on one bona fide source, we investigate what would be the resulting BPCER on the other data sources. To show that, Table 6 listed the BPCER values for all other bona fide images of the other data sources (second column) when fixing the decision threshold so the BPCER of the data source (in the first column) is at 0.2. Along each row, the lowest BPCER (with the most effective quality measure as an MAD score) is marked in bold. Under most settings, MagFace, the inverted RankIQ, BRISQUE, inverted CNNIQA, and dipIQ show low BPCER on other bona fide sources when counting the occurrences of minimum BPCER in bold across settings. We can see from Table 6 that despite the change in the bona fide source, thresholding the quality on

one source does not lead to extremely high BPCER on other data sources in most cases. If we neglected the unrealistic MorGAN bona fide, the FIQA methods perform significantly better in detecting bona fide than IQA methods.

How detectable are the different morphing attacks by quality when fixing the decision threshold on diverse bona fide sources (i.e. BPCER)?

To answer this, we fix the BPCER, each time on the bona fide of one of the five data sources and calculate the APCER of each of the attack types in each of the morphing datasets. Table 7 shows the APCER of the morphing attacks when fixing the BPCER of the individual bona fide scores at 0.2 for the bona fide source in the first column. Considering the first row, the experimental setting is to fix, for example, the decision threshold using a quality measure for the bona fide distribution of the FRLL-Morphs at BPCER = 0.2 and use this decision threshold to determine the APCER value of AMSL, FaceMorpher, OpenCV, StyleGAN 2, and WebMorph attacks in the FRLL-Morphs. The APCER value is not only determined for all morphing attack types within the same morphing dataset but also across all other attacks in the other morphing datasets as well. This is targeted at studying the effectiveness of using IQ and FIQ measures for MAD on unknown morphing attacks and bona fide sources.

The bold number in Table 7 shows the lowest APCER within the same morphing attacks and morphing dataset across different quality measures that are used as MAD scores and the bona fide source is used to fix the BPCER threshold. MagFace shows the lowest APCER for StyleGAN 2 generated morphs both in the FRLL-Morphs (APCER = 0.0123), FERET-Morphs (APCER = 0.0813), and the FRGC-Morphs (APCER = 0.0861). In comparison to StyleGAN 2 attacks, the APCER produced by MagFace is higher for landmark-based attacks, such as FaceMorpher, OpenCV, and WebMorph. On the other hand, the inverted CNNIQA shows low APCER values on landmark-based attacks, both within the same morphing dataset or in other morphing datasets, and this is more significant for the FRLL-Morphs and FERET-Morphs. Looking at the third row showing APCER results for FRGC-Morphs, the inverted dipIQ outperforms the inverted CNNIQA in terms of lower APCER for landmark-based morphing attacks.

For detecting the morphing attacks in the LMA-DRD datasets, MagFace outperforms all other IQA and FIQA measures by a large margin when the BPCER threshold is set on the same bona fide source. In both cases of LMA-DRD (D) and LMA-DRD (PS) datasets, MagFace and the inverted RankIQ clearly outperform all other quality measures based on IQA. The reason could be the mitigated image artefacts after the re-digitisation process, rendering the FIQA methods more superior detection performance.

In the case of the MorGAN dataset, due to its relatively low image resolution, we focussed only on the intra-class performance and its APCER. The lowest APCER is again observed for MagFace with 0.61 and 0.66 both for landmark-based OpenCV method and MorGAN-based morphing attacks, respectively, which is still very high.

**TABLE 6** BPCER of the morph datasets is given in the table, by fixing the decision threshold so that the BPCER of one morph dataset is at 0.2 (left-most column) and using this threshold to calculate the BPCER for other morphing datasets (second from left column). The bold numbers mark the best performing quality metric within each morphing dataset, which is determined across each row. One can see that despite the change in the bona fide source, thresholding the quality on one dataset does not lead to extremely high BPCER on other datasets in most cases

Dataset (BF)	dataset (BF)	MagFace	SERFIQ	RankIQ*	BRISQUE	CNNIQA*	dipIQ*
FRLM-morphs	FERET-morphs	0.3008	0.2788	0.1415	0.9766	0.2887	<b>0.0248</b>
	FRGC-morphs	0.3098	0.3249	0.1042	0.5513	0.6227	<b>0.0736</b>
	LMA-DRD (D)	<b>0.1689</b>	0.2098	0.0027	0.4687	0.8142	0.2207
	LMA-DRD (PS)	0.2039	0.2039	<b>0.0138</b>	0.3416	0.7438	0.2259
FERET-morphs	MorGAN	1.0000	1.0000	0.0010	0.9743	0.1316	<b>0.0274</b>
	FRLM-morphs	0.0588	<b>0.1716</b>	0.3578	0.0000	0.0882	0.5098
	FRGC-morphs	0.1876	0.2346	0.1734	<b>0.0224</b>	0.5463	0.2993
	LMA-DRD (D)	0.0981	0.1444	0.0245	<b>0.0054</b>	0.7486	0.5668
FRGC-morphs	LMA-DRD (PS)	0.1350	0.1460	0.0303	<b>0.0028</b>	0.6749	0.5675
	MorGAN	1.0000	1.0000	<b>0.0023</b>	0.1346	0.0986	0.1764
	FRLM-morphs	<b>0.0686</b>	0.1716	0.4020	0.0000	0.0049	0.3775
	FERET-morphs	0.2088	0.1635	0.2173	0.7481	<b>0.0113</b>	0.1203
LMA-DRD (D)	LMA-DRD (D)	0.1035	0.1144	<b>0.0245</b>	0.1090	0.3224	0.4142
	LMA-DRD (PS)	0.1433	0.1157	<b>0.0386</b>	0.0413	0.3691	0.4132
	MorGAN	1.0000	1.0000	<b>0.0027</b>	0.7348	0.0237	0.1015
	FRLM-morphs	0.2500	0.1961	0.7010	0.0049	<b>0.0000</b>	0.1569
LMA-DRD (PS)	FERET-morphs	0.3418	0.2562	0.3857	0.8662	<b>0.0014</b>	0.0191
	FRGC-morphs	0.3483	0.2952	0.4054	0.2760	0.0821	<b>0.0641</b>
	LMA-DRD (PS)	0.2314	0.1791	0.2066	<b>0.0826</b>	0.2011	0.1983
	MorGAN	1.0000	1.0000	0.0307	0.8474	<b>0.0107</b>	0.0204
MorGAN	FRLM-morphs	0.1814	0.1961	0.6961	0.0588	<b>0.0000</b>	0.1569
	FERET-morphs	0.2880	0.2725	0.3808	0.9519	<b>0.0014</b>	0.0198
	FRGC-morphs	0.2965	0.3142	0.3997	0.4263	<b>0.0815</b>	0.0654
	LMA-DRD (D)	<b>0.1553</b>	0.2071	0.1826	0.3842	0.1913	0.2071
MorGAN	MorGAN	1.0000	1.0000	0.0294	0.9509	<b>0.0104</b>	0.0214
	FRLM-morphs	<b>0.0000</b>	0.0000	0.8971	0.0000	0.4706	0.5490
	FERET-morphs	<b>0.0000</b>	0.0000	0.6100	0.2689	0.4416	0.2180
	FRGC-morphs	<b>0.0000</b>	0.0003	0.6681	0.0325	0.7483	0.3331
MorGAN	LMA-DRD (D)	<b>0.0000</b>	0.0000	0.5967	0.0191	0.8852	0.5913
	LMA-DRD (PS)	<b>0.0000</b>	0.0028	0.5840	0.0028	0.8347	0.6171

5.2.3 | Given an overall MAD performance measure, how well do quality measures perform as MADs across datasets and attacks?

To answer that, we investigate the ACER values scored on different datasets when the BPCER is fixed on each bona fide source individually. Table 8 shows the ACER values, which is the balanced error between BPCER and APCER. It shows a clearer trade-off between bona fide and attack detection. A

smaller ACER indicates the overall high performance of using this quality measure for the MAD task. This value is calculated when fixing the decision threshold such that the BPCER for the bona fide source shown in the first column is at 0.2. This threshold is then used to calculate the ACER, by adding the BPCER and APCER scores at this decision threshold and dividing by 2.

The ACER values in Table 8 provide supporting findings to those extracted for the APCER in Table 7 and BPCER

**TABLE 7** The APCER value of the morphing attacks is given in the table at the decision threshold that leads to the BPCER of 0.2 on the dataset in the left-most column. This table aims at showing the morphing detectability using the quality and utility measures within and across morphing attacks and morphing datasets. The bold number shows the lowest APCER error within the same morphing attacks and morphing dataset across different quality measures. Comparison within and across morphing datasets intend to show the MAD performance on known and unknown attacks and its generalisability on unseen morphing images

Dataset	Quality metric	FRLI-Morphs				FERET-Morphs				FRGC-Morphs				LMD-DRD (PS)		LMD-DRD (D)		MorGAN	
		Face		StyleGAN		FaceMorpher		OpenCV 2		FaceMorpher		OpenCV 2		OpenCV (PS)		OpenCV (D)		OpenCV MorGAN	
		AMSL	Morpher	OpenCV	2	WebMorph	FaceMorpher	OpenCV	2	StyleGAN	OpenCV	2	FaceMorpher	OpenCV	2	StyleGAN	OpenCV	2	OpenCV
FRLI-Morphs	MagFace	0.4207	0.3560	0.3399	<b>0.0123</b>	0.3530	0.1871	0.1871	<b>0.0359</b>	0.2925	0.3205	<b>0.0373</b>	0.3841	0.3127	<b>0.0000</b>	<b>0.0000</b>	0.0000	<b>0.0000</b>	
	SERFIQ	0.9007	0.8535	0.8559	0.2946	0.8460	0.6919	0.7108	0.4234	0.7376	0.7552	0.3268	0.6449	0.6182	<b>0.0000</b>	<b>0.0000</b>	0.0000	<b>0.0000</b>	
	RankIQ*	0.5857	0.5303	0.5299	0.7995	0.6626	0.6560	0.6635	0.7713	0.5093	0.5104	0.8195	0.9819	0.9891	0.9990	1.0000	0.9990	1.0000	
	BRISQUE	0.0239	0.0417	0.2432	0.2938	0.0074	0.1796	0.3743	0.1399	0.2355	0.4118	0.3714	0.2572	0.5927	0.0060	0.1360	0.0060	0.1360	
FERET-Morphs	CNNIQA*	<b>0.0138</b>	<b>0.0008</b>	<b>0.0049</b>	0.7512	<b>0.0008</b>	<b>0.0114</b>	<b>0.0170</b>	0.2760	<b>0.0000</b>	<b>0.0041</b>	0.1058	<b>0.0580</b>	<b>0.1091</b>	0.7750	0.8680	0.7750	0.8680	
	dipIQ*	0.9499	0.5025	0.0958	0.4280	0.9844	0.2684	0.3327	0.5331	0.3506	0.6131	0.6359	0.8007	0.5891	0.9790	0.9640	0.9790	0.9640	
	MagFace	0.5756	0.5442	0.5037	<b>0.0745</b>	0.5389	0.3214	0.3478	<b>0.0813</b>	0.4699	0.4824	<b>0.0944</b>	0.5362	0.4836	<b>0.0000</b>	<b>0.0000</b>	<b>0.0000</b>	<b>0.0000</b>	
	SERFIQ	0.9490	0.9157	0.9353	0.4935	0.9132	0.7996	0.8147	0.5633	0.8506	0.8351	0.5083	0.7464	0.7527	<b>0.0000</b>	<b>0.0000</b>	<b>0.0000</b>	<b>0.0000</b>	
FRGC-Morphs	RankIQ*	0.3756	0.3355	0.3235	0.5532	0.4201	0.5803	0.5784	0.6881	0.3371	0.3496	0.6649	0.9457	0.9491	0.9970	1.0000	0.9970	1.0000	
	BRISQUE	1.0000	1.0000	1.0000	1.0000	1.0000	0.9924	1.0000	0.9773	1.0000	1.0000	1.0000	0.9855	1.0000	0.8010	0.9790	0.8010	0.9790	
	CNNIQA*	<b>0.0925</b>	<b>0.0074</b>	<b>0.0139</b>	0.9026	<b>0.0442</b>	<b>0.0227</b>	<b>0.0491</b>	0.4291	<b>0.0042</b>	<b>0.0197</b>	0.2718	<b>0.1051</b>	<b>0.1636</b>	0.8320	0.9090	0.8320	0.9090	
	dipIQ*	0.0947	0.0106	0.0123	0.1841	0.2940	0.0378	0.1153	0.3138	0.0882	0.2956	0.3506	0.5072	0.3055	0.9020	0.7820	0.9020	0.7820	
LMD-DRD (PS)	MagFace	0.5549	0.5262	0.4775	<b>0.0655</b>	0.5184	0.3100	0.3308	<b>0.0699</b>	0.4512	0.4627	<b>0.0861</b>	<b>0.5145</b>	0.4655	<b>0.0000</b>	<b>0.0000</b>	<b>0.0000</b>	<b>0.0000</b>	
	SERFIQ	0.9678	0.9370	0.9533	0.6088	0.9369	0.8488	0.8582	0.6484	0.8859	0.8734	0.5695	0.8116	0.7818	<b>0.0000</b>	<b>0.0000</b>	<b>0.0000</b>	<b>0.0000</b>	
	RankIQ*	0.3232	0.2930	0.2785	0.4943	<b>0.3628</b>	0.5595	0.5614	0.6654	0.2956	<b>0.3060</b>	0.6307	0.9203	0.9345	0.9970	1.0000	0.9970	1.0000	
	BRISQUE	0.5968	0.7692	0.9853	0.9738	0.6830	0.7391	0.8866	0.6333	0.9336	0.9813	0.9741	0.7428	0.9636	0.1700	0.6340	0.1700	0.6340	
LMD-DRD (D)	CNNIQA*	0.9489	0.8609	0.8509	0.9992	0.9500	0.5833	0.5614	0.8507	0.5909	0.8154	0.9305	0.5942	0.5164	0.9660	0.9880	0.9660	0.9880	
	dipIQ*	<b>0.2786</b>	<b>0.0458</b>	<b>0.0221</b>	0.2651	0.6552	<b>0.0756</b>	<b>0.1720</b>	0.3951	<b>0.1639</b>	0.4108	0.4492	0.5942	<b>0.3709</b>	0.9460	0.8620	0.9460	0.8620	
	SERFIQ	0.9172	0.8707	0.8812	0.3576	0.8731	0.7259	0.7316	0.4707	0.7770	0.7790	0.3859	0.6703	0.6800	<b>0.0000</b>	<b>0.0000</b>	<b>0.0000</b>	<b>0.0000</b>	
	RankIQ*	<b>0.0887</b>	<b>0.1097</b>	<b>0.1073</b>	0.1489	<b>0.1237</b>	0.3989	0.3913	0.4594	<b>0.1110</b>	<b>0.1141</b>	0.3226	0.6486	0.6727	0.9530	0.9900	0.9530	0.9900	
LMD-DRD (PS)	BRISQUE	0.2864	0.5025	0.9009	0.8912	0.3292	0.5936	0.8147	0.4934	0.8268	0.9263	0.8932	0.6341	0.9236	0.0820	0.5010	0.0820	0.5010	
	CNNIQA*	0.9995	0.9812	0.9631	1.0000	0.9918	0.8030	0.8034	0.9338	0.9232	0.9678	0.9834	0.7283	0.7673	0.9850	0.9970	0.9850	0.9970	
	dipIQ*	0.9651	0.5908	0.1204	0.4509	0.9853	0.3081	0.3516	0.5520	0.3869	0.6463	0.6535	0.8225	0.6182	0.9800	0.9710	0.9800	0.9710	
	MagFace	0.4428	0.3781	0.3612	<b>0.0188</b>	0.3767	<b>0.1947</b>	<b>0.2023</b>	<b>0.0378</b>	<b>0.3133</b>	0.3402	<b>0.0415</b>	0.4130	<b>0.3200</b>	<b>0.0000</b>	<b>0.0000</b>	<b>0.0000</b>	<b>0.0000</b>	
SERFIQ	0.9053	0.8592	0.8632	0.3142	0.8509	0.7070	0.7164	0.4367	0.7500	0.7604	0.3423	0.6486	0.6255	<b>0.0000</b>	<b>0.0000</b>	<b>0.0000</b>	<b>0.0000</b>		

TABLE 7 (Continued)

Dataset	Quality metric	FRLI-Morphs				FERET-Morphs				FRGC-Morphs				LMD-DRD (D)		LMD-DRD (PS)		MorGAN			
		Face		StyleGAN 2		WebMorph		StyleGAN 2		FaceMorpher		OpenCV 2		StyleGAN 2		OpenCV (D)		OpenCV (PS)		MorGAN	
		AMSL	Morpher	OpenCV		OpenCV		FaceMorpher	OpenCV	OpenCV		FaceMorpher	OpenCV	OpenCV		OpenCV (D)	OpenCV (PS)	OpenCV	MorGAN	OpenCV	MorGAN
MorGAN	RankIQ*	0.0929	<b>0.1105</b>	<b>0.1089</b>	0.1555	0.1294	0.4026	0.3970	0.4631	0.1131	<b>0.1193</b>	0.3330	0.5558	0.6873	0.5530	0.9900					
	BRISQUE	<b>0.0782</b>	0.1481	0.4898	0.5589	<b>0.0450</b>	0.3043	0.5425	0.2514	0.4346	0.6504	0.6058	<b>0.3732</b>	0.7491	0.0170	0.2320					
	CNNIQA*	0.9995	0.9812	0.9631	1.0000	0.9918	0.8030	0.8034	0.9338	0.9263	0.9678	0.9834	0.7319	0.7673	0.9850	0.9970					
	dipIQ*	0.9646	0.5810	0.1171	0.4493	0.9844	0.3043	0.3516	0.5520	0.3828	0.6421	0.6515	0.8225	0.6145	0.9800	0.9710					
	MagFace	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	<b>0.6100</b>	<b>0.6620</b>					
	SERFIQ	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.7700	0.6800					
	RankIQ*	0.0211	0.0417	0.0467	<b>0.0336</b>	<b>0.0590</b>	0.2136	0.2079	0.2344	0.0207	0.0322	0.0871	0.2681	0.3055	0.7000	0.8660					
	BRISQU	1.0000	1.0000	1.0000	1.0000	1.0000	0.9924	0.9981	0.9622	1.0000	1.0000	1.0000	0.9819	1.0000	0.7290	0.9630					
	CNNIQA*	<b>0.0000</b>	<b>0.0000</b>	<b>0.0008</b>	0.3674	0.0000	<b>0.0000</b>	<b>0.0095</b>	<b>0.1323</b>	<b>0.0000</b>	<b>0.0000</b>	<b>0.0197</b>	<b>0.0362</b>	<b>0.0764</b>	0.6890	0.7850					
	dipIQ*	0.0699	0.0065	0.0082	0.1702	0.2506	0.0340	0.1040	0.2987	0.0768	0.2604	0.3133	0.4601	0.2655	0.8850	0.7340					

values in Table 6. The bold number shows the lowest ACER values determined when fixing the BPCER of the individual bona fide source in the first column at 0.2. MagFace as the detection metric performs more stable across all different kinds of morphing attacks and morphing datasets. For FRLI-Morphs, FERET-Morphs, and FRGC-Morphs, MagFace shows, among other quality measures the best ACER for detecting morphing attacks such as StyleGAN 2 and also the attacks in LMA-DRD. The inverted CNN and dipIQ share the success for detecting landmark-based morphing attacks with minimum ACER value in the individual experimental settings. For the re-digitised LMA-DRD dataset, BRISQUE shows good error performance in terms of minimum ACER.

The right-most column in Table 8 shows the *mean* average ACER value across all morphing attacks and morphing datasets for each IQA and FIQA measure. We also marked all *mean* ACER values below 0.3 in bold for better visibility. Based on this averaged ACER value, we observe that several quality/utility measures show stable ACER in the MAD task across different morphing attacks and morphing datasets. Mostly MagFace, but also the inverted CNNIQA, and the inverted dipIQ all perform equally well under specific experimental conditions. While FIQA measures work well on StyleGAN 2 generated morphing attacks, IQA measures work well in detecting facial landmark-based morphing attacks. The good performance of these measures is reflected in the low ACER values. Thus the quality measures of MagFace and CNNIQA can be used as unsupervised MAD with the overall performance of ACER below 0.30 under the diverse and unseen bona fide sources and attack variations. This relatively high MAD performance by MagFace and CNNIQA might be explained by their design concept. The MagFace quality is the magnitude of the sample embedding using the network trained as proposed in Ref. [14]. This is based on training the FR model using a loss that adapts the penalty margin loss based on this magnitude and thus links the closeness of a sample to its class centre to the unnormalised embedding magnitude. A morphed face image is designed to be close to multiple classes (identities) and thus is typically not expected to be extremely close to an individual class. The CNNIQA quality is based on processing image patches. This property is proven in Ref. [13] to be extremely sensitive to local distortions, which makes it a perfect choice for local quality estimation. This corresponds to being sensitive to the local artefacts caused by the different morphing processes as discussed and shown in Ref. [8]. One can take advantage of the different properties of image IQA and FIQA methods to build a possibly stronger MAD based on quality measures. This can be performed by quality score-level information fusion approach [86], which, after score normalisation [87], can give relative weights that correspond to the estimated performance of each of the fused methods [81, 88].

To sum up, our investigation on the detectability of morphing attacks using the quality measures shows that FIQA and IQA metrics can be used as unsupervised attack detection measures under different attack scenarios. While MagFace from FIQA performs relatively better on GAN-based



TABLE 8 (Continued)

Dataset	Quality metric	FRLL-Morphs			FERET-Morphs			FRGC-Morphs			LMD-DRD (D)		LMD-DRD (PS)		MorGAN		
		AMSL	FaceMorpher	OpenCV	StyleGAN 2	WebMorph	FaceMorpher	OpenCV	StyleGAN 2	FaceMorpher	OpenCV	OpenCV (D)	OpenCV (PS)	OpenCV (D)	OpenCV (PS)	OpenCV	MorGAN Mean
MorGAN	MagFace	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	<b>0.4050</b>	<b>0.4310</b>	0.4890
	SERFIQ	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	0.5002	0.5000	0.5014	0.5014	0.5014	0.4847	0.4397	0.4950
	RankIQ*	0.4591	0.4694	0.4719	0.4653	0.4780	0.4118	0.4090	0.4222	0.3444	0.3501	<b>0.4324</b>	0.4447	0.4447	0.4500	0.5330	0.4345
	BRISQUE	0.5000	0.5000	0.5000	0.5000	0.5000	0.6307	0.6335	0.6156	0.5163	0.5163	0.5005	0.5014	0.5014	0.4645	0.5815	0.5317
	CNNIQA*	<b>0.2353</b>	<b>0.2353</b>	<b>0.2357</b>	0.4190	<b>0.2353</b>	0.2208	0.2255	0.2870	0.3742	0.3742	0.4607	0.4555	0.4555	0.4456	0.4936	0.3387
	dipIQ*	0.3095	0.2778	0.2786	<b>0.3596</b>	0.3998	<b>0.1260</b>	<b>0.1610</b>	<b>0.2583</b>	<b>0.2049</b>	<b>0.2967</b>	<b>0.3232</b>	<b>0.4413</b>	<b>0.4413</b>	0.5427	0.4672	<b>0.3314</b>

morphing attacks, the inverse CNNIQA from IQA measure performs relatively better on landmark-based morphing attacks.

### 5.3 | Take-home messages

In this section, we summarise and highlight the most important outcomes discussed in detail in Section 5:

- All face morphing techniques do introduce a change in FIQ and IQ measures of face images in comparison to their corresponding bona fide samples. Such a shift was larger and more consistent in some FIQ/IQ methods than in others, for example, MagFace, CNNIQA, and dipIQ showed relatively stronger quality score shifts. This shift directly corresponds to a quality separability between bona fide and attacks.
- The proposed use of FIQ and IQ measures as unsupervised MAD methods produce a stable and generalisable MAD performance for some of the quality measures. We noticed that setting an MAD decision threshold can produce relatively generalisable MAD decisions across bona fide sources and morphing techniques. This is especially the case for MagFace, CNNIQA, and dipIQ, as expected given the stable effect of different morphing techniques on their quality scores.
- Following the above-mentioned observations, we showed that using quality measures, such as MagFace and CNNIQA, as unsupervised MAD results in ACER values below 30% on a diverse set of unknown data sources and morphing techniques.

## 6 | CONCLUSION

Given the complexity of MAD and the threat the morphing operation posed on the automatic face recognition system, we studied the effect of morphing on image quality and utility. We found a general quality shift in morphed images compared to bona fide images by using these investigated IQA and FIQA measures. However, these quality measures behave differently on landmark-based and GAN-based morphing attacks, yet to different degrees.

Based on the observed separability in the quality measures between the bona fide and morph samples, we theorise that such measures can be used as unsupervised MADs, where the quality score can act as an MAD detection score. Our analyses were successful in showing that using these quality measures to differentiate between morphing attacks and bona fide samples, quality measures such as MagFace and inverted CNNIQA can consistently lead to ACER values below 0.30 on completely unknown bona fide sources and attack methods. Based on that, we conclude that these quality measures represent useful decision measures for practitioners designing applications for the MAD task as a stand-alone MAD indicator, or in future work, as a supporting measure for MAD decisions.

## ACKNOWLEDGEMENTS

This research work has been funded by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE.

## CONFLICT OF INTEREST

None.

## DATA AVAILABILITY STATEMENT

Data sharing is not applicable to this article as no datasets were generated or analysed during the current study.

## ORCID

*Biying Fu*  <https://orcid.org/0000-0003-4848-1256>

*Naser Damer*  <https://orcid.org/0000-0001-7910-7895>

## REFERENCES

- Boutros, F., et al.: Elasticface: elastic margin loss for deep face recognition. In: IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops. Computer Vision Foundation / IEEE (2022)
- Deng, J., et al.: Arcface: additive angular margin loss for deep face recognition. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019, pp. 4690–4699. Computer Vision Foundation / IEEE (2019)
- Boutros, F., et al.: Mixfacenets: extremely efficient face recognition networks. In: International IEEE Joint Conference on Biometrics, IJCB 2021, pp. 1–8. IEEE (2021)
- Boutros, F., et al.: Pocketnet: extreme lightweight face recognition network using neural architecture search and multistep knowledge distillation. IEEE Access 10, 46823–46833 (2022). <https://doi.org/10.1109/access.2022.3170561>
- Ferrara, M., Franco, A., Maltoni, D.: The magic passport. In: IEEE International Joint Conference on Biometrics, Clearwater, IJCB 2014, FL, USA, September 29 - October 2, 2014, pp. 1–7. IEEE (2014)
- Damer, N., Boutros, F., Alexandra Mosegui S., Kirchbuchner, F., Kuijper, A.: Realistic dreams: cascaded enhancement of gan-generated images with an example in face morphing attacks. In: 10th IEEE International Conference on Biometrics Theory, Applications and Systems, BTAS 2019, pp. 1–10. IEEE (2019)
- Zhang, H., et al.: MIPGAN - generating strong and high quality morphing attacks using identity prior driven GAN. IEEE Trans. Biom. Behav. Identity Sci. 3(3), 365–383 (2021). <https://doi.org/10.1109/tbiom.2021.3072349>
- Damer, N., et al.: ReGenMorph: visibly realistic GAN generated face morphing attacks by attack re-generation. In: Advances in Visual Computing - 16th International Symposium, ISVC 2021, Virtual Event, October 4-6, 2021, Proceedings, Part I, Volume 13017 of Lecture Notes in Computer Science, pp. 251–264. Springer (2021)
- Neubert, T., et al.: Extended StirTrace benchmarking of biometric and forensic qualities of morphed face images. IET Biom. 7(4), 325–332 (2018). <https://doi.org/10.1049/iet-bmt.2017.0147>
- Debiasi, L., Damer, N., Alexandra Mosegui S., Rathgeb, C., Ulrich, S., Busch, C., Kirchbuchner, F., Uhl, A.: On the detection of gan-based face morphs using established morph detectors. In: Image Analysis and Processing - ICIAP 2019 - 20th International Conference, Trento, Italy, September 9-13, 2019, Proceedings, Part II, Volume 11752 of Lecture Notes in Computer Science, pp. 345–356. Springer (2019)
- Fu, B., et al.: The effect of face morphing on face image quality. In: Proceedings of the 20th International Conference of the Biometrics Special Interest Group, BIOSIG 2021, Digital Conference, September 15-17, 2021, Volume P-315 of LNI, pp. 173–180. Gesellschaft für Informatik e.V. (2021)
- Mallik, S.: Face Morph Using Opencv — C++/python (2016)
- Le Kang, et al.: Doermann. Convolutional neural networks for no-reference image quality assessment. In: 2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014, pp. 1733–1740. IEEE Computer Society (2014)
- Meng, Q., et al.: MagFace: a universal representation for face recognition and quality assessment. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, Virtual, June 19-25, 2021, pp. 14225–14234. Computer Vision Foundation / IEEE (2021)
- Ferrara, M., Franco, A., Maltoni, D.: The magic passport. In: IEEE International Joint Conference on Biometrics, pp. 1–7. IEEE (2014)
- Moorthy, A.K., Bovik, A.C.: A two-step framework for constructing blind image quality indices. IEEE Signal Process. Lett. 17(5), 513–516 (2010). <https://doi.org/10.1109/lsp.2010.2043888>
- Shao, W., Mou, X.: No-reference image quality assessment based on edge pattern feature in the spatial domain. IEEE Access 9, 133170–133184 (2021). <https://doi.org/10.1109/access.2021.3112741>
- Liu, L., et al.: Blind image quality assessment by relative gradient statistics and adaboosting neural network. Signal Process. Image Commun. 40, 1–15 (2016). <https://doi.org/10.1016/j.image.2015.10.005>
- Liu, L., et al.: No-reference image quality assessment based on spatial and spectral entropies. Signal Process. Image Commun. 29(8), 856–863 (2014). <https://doi.org/10.1016/j.image.2014.06.006>
- Pankaj Shivdayal Wasnik, Raja, K.B., Raghavendra, R., Busch, C.: Assessing face image quality for smartphone based face recognition system. In: 5th International Workshop on Biometrics and Forensics, IWBF 2017, Coventry, United Kingdom, April 4-5, 2017, pp. 1–6. IEEE (2017)
- Gao, X., et al.: Standardization of face image sample quality. In: Advances in Biometrics, International Conference, ICB 2007, Seoul, Korea, August 27-29, 2007, Proceedings, Volume 4642 of Lecture Notes in Computer Science, pp. 242–251. Springer (2007)
- Niranjan, D.N., Karam, L.J.: A no-reference image blur metric based on the cumulative probability of blur detection (cpbd). IEEE Trans. Image Process. 20(9), 2678–2683 (2011). <https://doi.org/10.1109/tip.2011.2131660>
- Shirvaikar, M.V.: An optimal measure for camera focus and exposure. In: Thirty-Sixth Southeastern Symposium on System Theory, 2004. Proceedings of the, pp. 472–475. IEEE (2004)
- Matkovic, K., et al.: Global contrast factor - a new approach to image contrast. In: 1st Eurographics Workshop on Computational Aesthetics in Graphics, Visualization, and Imaging, CAe 2005, Girona, Spain, May 18-20, 2005, Proceedings, pp. 159–167. Eurographics Association (2005)
- Fatema Tuz Z., Andrei Dmitri G., Zatarain Duran, O., Marina, L., Gavrilova: A linear regression model for estimating facial image quality. In: 16th IEEE International Conference on Cognitive Informatics & Cognitive Computing, ICCI\*CC 2017, Oxford, United Kingdom, July 26-28, 2017, pp. 130–138. IEEE Computer Society (2017)
- Turaga, D.S., Chen, Y., Caviedes, J.E.: No reference PSNR estimation for compressed pictures. Signal Process. Image Commun. 19(2), 173–184 (2004). <https://doi.org/10.1016/j.image.2003.09.001>
- Dosselmann, R., XueYang, D.: A comprehensive assessment of the structural similarity index. Signal Image Video Process 5(1), 81–91 (2011). <https://doi.org/10.1007/s11760-009-0144-1>
- Terhörst, P., Jan Niklas K., Damer, N., Kirchbuchner, F., Kuijper, A.: Face quality estimation and its correlation to demographic and non-demographic bias in face recognition. In: 2020 IEEE International Joint Conference on Biometrics, IJCB 2020, Houston, TX, USA, September 28 - October 1, 2020, pp. 1–11. IEEE (2020)
- Hernandez-Ortega, J., et al.: Faceqnet: quality assessment for face recognition based on deep learning. In: 2019 International Conference on Biometrics, ICB 2019, Crete, Greece, June 4-7, 2019, pp. 1–8. IEEE (2019)
- Chen, J., et al.: Face image quality assessment based on learning to rank. IEEE Signal Process. Lett. 22(1), 90–94 (2015). <https://doi.org/10.1109/lsp.2014.2347419>

31. Venkatanath, N., et al.: Blind image quality evaluation using perception based features. In: Twenty First National Conference on Communications, NCC 2015, Mumbai, India, February 27 - March 1, 2015, pp. 1–6. IEEE (2015)
32. Mittal, A., Soundararajan, R., Alan Bovik, C.: Making a completely blind image quality analyzer. *IEEE Signal Process. Lett.* 20(3), 209–212 (2013). <https://doi.org/10.1109/lsp.2012.2227726>
33. Bosse, S., et al.: Deep neural networks for no-reference and full-reference image quality assessment. *IEEE Trans. Image Process.* 27(1), 206–219 (2018). <https://doi.org/10.1109/tip.2017.2760518>
34. Ma, K., et al.: End-to-end blind image quality assessment using deep neural networks. *IEEE Trans. Image Process.* 27(3), 1202–1213 (2018). <https://doi.org/10.1109/tip.2017.2774045>
35. Liu, X., et al.: Learning from rankings for no-reference image quality assessment. In: IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22–29, 2017, pp. 1040–1049. IEEE Computer Society (2017)
36. Ma, K., et al.: dipIQ: blind image quality assessment by learning-to-rank discriminable image pairs. *IEEE Trans. Image Process.* 26(8), 3951–3964 (2017). <https://doi.org/10.1109/tip.2017.2708503>
37. Venkatesh, S., et al.: Face morphing attack generation and detection: a comprehensive survey. *IEEE Transactions on Technology and Society* 2(3), 128–145 (2021). <https://doi.org/10.1109/tts.2021.3066254>
38. Aghdaie, P., et al.: Detection of morphed face images using discriminative wavelet sub-bands. In: 9th IEEE International Workshop on Biometrics and Forensics, IWBF 2021, pp. 1–6. IEEE (2021)
39. Damer, N., et al.: PW-MAD: pixel-wise supervision for generalized face morphing attack detection. In: Advances in Visual Computing - 16th International Symposium, ISVC 2021, Virtual Event, October 4–6, 2021, Proceedings, Part I, Volume 13017 of Lecture Notes in Computer Science, pp. 291–304. Springer (2021)
40. Damer, N., Zienert, S., Wainakh, Y., Alexandra Mosegui S., Kirchbuchner, F., Kuijper, A.: A multi-detector solution towards an accurate and generalized detection of face morphing attacks. In: 22th International Conference on Information Fusion, FUSION 2019, Ottawa, ON, Canada, July 2–5, 2019, pp. 1–8. IEEE (2019)
41. Raghavendra, R., et al.: Face morphing versus face averaging: vulnerability and detection. In: 2017 IEEE International Joint Conference on Biometrics, IJCB 2017, Denver, CO, USA, October 1–4, 2017, pp. 555–563. IEEE (2017)
42. Raghavendra, R., et al.: Transferable deep-cnn features for detecting digital and print-scanned morphed face images. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2017, Honolulu, HI, USA, July 21–26, 2017, pp. 1822–1830. IEEE Computer Society (2017)
43. Damer, N., et al.: Detecting face morphing attacks by analyzing the directed distances of facial landmarks shifts. In: Pattern Recognition - 40th German Conference, GCPR 2018, Stuttgart, Germany, October 9–12, 2018, Proceedings, Volume 11269 of Lecture Notes in Computer Science, pp. 518–534. Springer (2018)
44. Damer, N., Alexandra Mosegui S., Zienert, S., Wainakh, Y., Terhörst, P., Kirchbuchner, F., Kuijper, A.: To detect or not to detect: the right faces to morph. In: 2019 International Conference on Biometrics, ICB 2019, Crete, Greece, June 4–7, 2019, pp. 1–8. IEEE (2019)
45. Ulrich, S., et al.: Deep face representations for differential morphing attack detection. *IEEE Trans. Inf. Forensics Secur.* 15, 3625–3639 (2020). <https://doi.org/10.1109/tifs.2020.2994750>
46. Soleymani, S., et al.: Differential morphed face detection using deep siamese networks. In: Pattern Recognition. ICPR International Workshops and Challenges - Virtual Event, January 10–15, 2021, Proceedings, Part VI, Volume 12666 of Lecture Notes in Computer Science, pp. 560–572. Springer (2020)
47. Raghavendra, R., Kiran, B., Raja, and Christoph Busch. Detecting morphed face images. In: 8th IEEE International Conference on Biometrics Theory, Applications and Systems, BTAS 2016, Niagara Falls, NY, USA, September 6–9, 2016, pp. 1–7. IEEE (2016)
48. Ulrich, S., et al.: On the vulnerability of face recognition systems towards morphed face attacks. In: 5th International Workshop on Biometrics and Forensics, IWBF 2017, Coventry, United Kingdom, April 4–5, 2017, pp. 1–6. IEEE (2017)
49. Damer, N., Alexandra Mosegui S., Braun, A., MorGAN, A.K.: Recognition vulnerability and attack detectability of face morphing attacks created by generative adversarial network. In: 9th IEEE International Conference on Biometrics Theory, Applications and Systems, BTAS 2018, Redondo Beach, CA, USA, October 22–25, 2018, pp. 1–10. IEEE (2018)
50. Ferrara, M., Franco, A., Maltoni, D.: Face morphing detection in the presence of printing/scanning and heterogeneous image sources. *IEE Biom.* 10(3), 290–303 (2021). <https://doi.org/10.1049/bme.2.12021>
51. Damer, N., César Augusto F.L., Fang, M., Spiller, N., Pham, M.Vu, Boutros, F.: Privacy-friendly synthetic data for the development of face morphing attack detectors. In: IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2022, New Orleans, Louisiana, USA, June 19–24, 2022. Computer Vision Foundation / IEEE (2022)
52. Damer, N., et al.: On the generalization of detecting face morphing attacks as anomalies: novelty vs. outlier detection. In: 10th IEEE International Conference on Biometrics Theory, Applications and Systems, BTAS 2019, Tampa, FL, USA, September 23–26, 2019, pp. 1–5. IEEE (2019)
53. Ulrich, S., Rathgeb, C., Busch, C.: Performance variation of morphed face image detection algorithms across different datasets. In: 2018 International Workshop on Biometrics and Forensics, IWBF 2018, Sassari, Italy, June 7–8, 2018, pp. 1–6. IEEE (2018)
54. Makrushin, A., Neubert, T., Jana, D.: Automatic generation and detection of visually faultless facial morphs. In: Proceedings of the 12th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2017) - Volume 6: VISAPP, Porto, Portugal, February 27 - March 1, 2017, pp. 39–50. SciTePress (2017)
55. Boutros, F., et al.: CR-FIQA: face image quality assessment by learning sample relative classifiability. *CoRR*, abs/2021.06592 (2021)
56. Wayman, J.L., et al.: Christoph busch, nicolas delvaux, and andrey zudenkov. Vocabulary harmonisation for biometrics: the development of ISO/IEC 2382 part 37, 2014
57. Fu, B., Kirchbuchner, F., Damer, N.: The effect of wearing a face mask on face image quality. In: 16th IEEE International Conference on Automatic Face and Gesture Recognition, FG 2021, Jodhpur, India, December 15–18, 2021, pp. 1–8. IEEE (2021)
58. Zhang, W., et al.: Blind image quality assessment using a deep bilinear convolutional neural network. *IEEE Trans. Circ. Syst. Video Technol.* 30(1), 36–47 (2020). <https://doi.org/10.1109/tcsvt.2018.2886771>
59. Ruderman, D.L.: The statistics of natural images. *Netw. Comput. Neural Syst.* 5(4), 517–548 (1994). [https://doi.org/10.1088/0954-898x\\_5\\_4\\_006](https://doi.org/10.1088/0954-898x_5_4_006)
60. Jia, D., et al.: Imagenet: a large-scale hierarchical image database. In: 2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20–25 June 2009, pp. 248–255. IEEE Computer Society (2009)
61. Zhang, W., et al.: Uncertainty-aware blind image quality assessment in the laboratory and wild. *IEEE Trans. Image Process.* 30, 3474–3486 (2021). <https://doi.org/10.1109/tip.2021.3061932>
62. Terhörst, P., Jan Niklas K., Damer, N., Kirchbuchner, F., Kuijper, A.: SER-FIQ: unsupervised estimation of face image quality based on stochastic embedding robustness. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13–19, 2020, pp. 5650–5659. Computer Vision Foundation / IEEE (2020)
63. Ou, Fu-Z., et al.: SDD-FIQA: unsupervised face image quality assessment with similarity distribution distance. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, Virtual, June 19–25, 2021, pp. 7670–7679. Computer Vision Foundation / IEEE (2021)
64. He, K., et al.: Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27–30, 2016, pp. 770–778. IEEE Computer Society (2016)

65. Fu, B., et al.: A deep insight into measuring face image utility with general and face-specific image quality metrics. In: IEEE/CVF Winter Conference on Applications of Computer Vision, WACV 2022, Waikoloa, HI, USA, January 3-8, 2022, pp. 1121–1130. IEEE (2022)
66. Fu, B., Damer, N.: Explainability of the implications of supervised and unsupervised face image quality estimations through activation map variation analyses in face recognition models. In: IEEE/CVF Winter Conference on Applications of Computer Vision Workshops, WACV - Workshops, Waikoloa, HI, USA, January 4-8, 2022, pp. 349–358. IEEE (2022)
67. DavisDlib-ml, E.K.: A machine learning toolkit. *J. Mach. Learn. Res.* 10, 1755–1758 (2009)
68. Sarkar, E., et al.: Vulnerability analysis of face morphing attacks from landmarks and generative adversarial networks. *CoRR*, abs/2005.05344 (2020)
69. Quek, A.: *Facemorpher* (2019)
70. Milborrow, S., Nicolls, F.: Active shape models with SIFT descriptors and MARS. In: *VISAPP 2014 - Proceedings of the 9th International Conference on Computer Vision Theory and Applications, Volume 2*, Lisbon, Portugal, 5-8 January, 2014, pp. 380–387. SciTePress (2014)
71. DeBruine, L.: *Debruine/Webmorph: Beta Release 2*. Jan 2018
72. DeBruine, L., Jones, B.: *Face Research Lab London Set, vol. 5* (2017)
73. Karras, T., et al.: Analyzing and improving the image quality of stylegan. In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pp. 8107–8116. Computer Vision Foundation / IEEE (2020)
74. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pp. 4401–4410. Computer Vision Foundation / IEEE (2019)
75. Venkatesh, S., et al.: Can GAN generated morphs threaten face recognition systems equally as landmark based morphs? - vulnerability and detection. In: *8th International Workshop on Biometrics and Forensics, IWBF 2020, Porto, Portugal, April 29-30, 2020*, pp. 1–6. IEEE (2020)
76. Jonathon Phillips, P., et al.: The FERET database and evaluation procedure for face-recognition algorithms. *Image Vis Comput.* 16(5), 295–306 (1998). [https://doi.org/10.1016/s0262-8856\(97\)00070-x](https://doi.org/10.1016/s0262-8856(97)00070-x)
77. Jonathon Phillips, P., et al.: Overview of the face recognition grand challenge. In: *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2005), 20-26 June 2005, San Diego, CA, USA*, pp. 947–954. IEEE Computer Society (2005)
78. Cao, Q., et al.: Vggface2: a dataset for recognising faces across pose and age. In: *13th IEEE International Conference on Automatic Face & Gesture Recognition, FG 2018, Xi'an, China, May 15-19, 2018*, pp. 67–74. IEEE Computer Society (2018)
79. Liu, Z., et al.: Deep learning face attributes in the wild. In: *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pp. 3730–3738. IEEE Computer Society (2015)
80. Zhang, K., et al.: Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Process. Lett.* 23(10), 1499–1503 (2016). <https://doi.org/10.1109/lsp.2016.2603342>
81. Damer, N., Alexander Opel, Alexander, N.: Biometric source weighting in multi-biometric fusion: towards a generalized and robust solution. In: *22nd European Signal Processing Conference, EUSIPCO 2014, Lisbon, Portugal, September 1-5, 2014*, pp. 1382–1386. IEEE (2014)
82. Standard for testing of presentation attack detection. *Int. Standard ISO/IEC30107-3*, (2016)
83. Information technology – Biometric sample quality – Part 1: Framework. *Int. Standard ISO/IEC 29794-1, (S2016)*
84. Damer, N., et al.: *Masked Face Recognition: Human versus Machine. IET Biom.* (2021)
85. Fang, M., et al.: Real masks and spoof faces: on the masked face presentation attack detection. *Pattern Recogn.* 123, 108398 (2022). <https://doi.org/10.1016/j.patcog.2021.108398>
86. Damer, N., Alexander Opel, Shahverdyan, A.: An overview on multi-biometric score-level fusion - verification and identification. In: *De Marsico, M., AnaFred, L.N. (eds.) ICPRAM 2013 - Proceedings of the 2nd International Conference on Pattern Recognition Applications and Methods, Barcelona, Spain, 15-18 February, 2013*, pp. 647–653. SciTePress (2013)
87. Damer, N., et al.: P-score: performance aligned normalization and an evaluation in score-level multi-biometric fusion. In: *26th European Signal Processing Conference, EUSIPCO 2018, Roma, Italy, September 3-7, 2018*, pp. 1402–1406. IEEE (2018)
88. Damer, N., Alexander Opel, Alexander, N.: CMC curve properties and biometric source weighting in multi-biometric score-level fusion. In: *17th International Conference on Information Fusion, FUSION 2014, Salamanca, Spain, July 7-10, 2014*, pp. 1–6. IEEE (2014)

**How to cite this article:** Fu, B., Damer, N.: Face morphing attacks and face image quality: The effect of morphing and the unsupervised attack detection by quality. *IET Biome.* 1–24 (2022). <https://doi.org/10.1049/bme2.12094>