

Received 20 October 2022, accepted 15 November 2022, date of publication 24 November 2022,  
date of current version 30 November 2022.

Digital Object Identifier 10.1109/ACCESS.2022.3224437

## RESEARCH ARTICLE

# Toward Robust Facial Authentication for Low-Power Edge-AI Consumer Devices

WANG YAO<sup>1</sup>, (Graduate Student Member, IEEE),  
VIKTOR VARKARAKIS<sup>2</sup>, (Graduate Student Member, IEEE), GABRIEL COSTACHE<sup>1,2</sup>,  
JOSEPH LEMLEY<sup>1,2</sup>, (Member, IEEE), AND PETER CORCORAN<sup>3</sup>, (Fellow, IEEE)

<sup>1</sup>College of Science and Engineering, University of Galway, H91 TK33 Galway, Ireland

<sup>2</sup>Xperi Corporation, H91 V0TX Galway, Ireland

<sup>3</sup>Institute of Electrical and Electronics Engineers, College of Science and Engineering, University of Galway, H91 TK33 Galway, Ireland

Corresponding author: Wang Yao (w.yao2@universityofgalway.ie)

This work was supported in part by the Irish Research Council Enterprise Partnership Ph.D. Scheme under Project EPSPG/2020/40, and in part by Xperi Corporation, Ireland.

**ABSTRACT** Robust authentication for low-power consumer devices without a keyboard remains a challenge. The recent availability of low-power neural accelerator hardware, combined with improvements in neural facial recognition algorithms provides enabling technology for low-power, on-device facial authentication. The present research work explores a number of approaches to test the robustness of a state-of-the-art facial recognition (FR) technique, Arcface for such end-to-end applications. As extreme lighting conditions and facial pose are the two more challenging scenarios for FR we focus on these. Due to the general lack of large-scale multiple-identity datasets, GAN-based re-lighting and pose techniques are used to explore the effects on FR performance. These results are further validated on the best available multi-identity datasets - MultiPIE and BIWI. The results show that FR is quite robust to pose variations up to 45-55 degrees, but the outcomes are not definitive for the tested lighting scenarios. For lighting, the tested GAN-based relighting augmentations show significant effects on FR robustness. However, the lighting scenarios from MultiPIE dataset - the best available public dataset - show some conflicting results. It is unclear if this is due to an incorrectly learned GAN relighting transformation or, alternatively, to mixed ambient/directional lighting scenes in the dataset. However, it is shown that the GAN-induced FR errors for extreme lighting conditions can be corrected by fine-tuning the FR network layers. The conclusions support the feasibility of implementing a robust authentication method for low-power consumer devices.

**INDEX TERMS** Directional lighting, face illumination, face recognition, face re-lighting method, FR evaluation, Internet of Things.

## I. INTRODUCTION

Human Face Recognition (FR) has been an active research field in computer vision since the early 1960s [1] and early Convolutional Neural Network (CNN) based approaches were in evidence before the end of that century [2]. Over the last two decades, FR has been well-studied in the literature with the most recent advances being driven by advances in CNN and deep learning [3], [4], [5], [6], [7]. In much of the literature the test samples for FR are assumed to be

The associate editor coordinating the review of this manuscript and approving it for publication was Michele Nappi.

normalized in terms of pose, facial expression and illumination to simplify the challenge of accurately distinguishing an individual identity among a very large population. But, as it is not always feasible to capture optimal facial samples, some studies have explored the effects of different factors on the accuracy of state-of-the-art (SoA) FR systems.

The main factors that affect FR include (i) pose [8], [9], [10], [11], (ii) illumination [12], [13], [14], (iii) facial expression [15], [16], [17], (iv) age [18], [19], (v) gender variations [20], [21], and (vi) facial occlusions [22], [23], [24].

In this work, the focus is on the latest end-to-end fully neural FR techniques [25] as these represent current SoA in

terms of accuracy and have the potential for implementation in the latest neural accelerators [26], [27]. The initial focus for the implementation of neural algorithms in embedded devices was on network optimizations such as parameter quantization and pruning, compressed convolutional filters and matrix factorizations [27]. However, the attention has recently shifted towards specialized neural topologies [28], [29] and ultra-low power realizations in hardware [30], [31]. Such optimizations enable SoA neural FR architectures to be implemented in a low-power consumer appliance, enabling a new generation of devices capable of identifying their owners, providing access control, and personalizing the device's responses and behavior.

Ideally, for optimal data security and user privacy, such devices should be capable to authenticate a user, without sending facial data to a network or cloud service. However, this introduces new challenges as such FR embodiment can no longer rely on pre-processing input facial samples to optimize power consumption. Thus, all image processing must be achieved in a fully neural implementation, requiring a neural end-to-end FR to be robust to factors such as pose and illumination. Here our goal is to determine the feasibility of modifying a high-accuracy SoA neural FR architecture to demonstrate robustness to uncompensated input image samples.

More specifically, this research studies the effects of directional illumination and larger pose facial samples in an end-to-end fully neural FR system. Key research questions include:

- 1) What metrics can be used to measure and compare FR at various directional lighting and different pose angles?
- 2) How robust are the latest FR algorithms to illumination and pose variations? In particular, can we quantify the effects of extreme lighting and pose variations on the robustness of FR?
- 3) Can a SoA FR algorithm be fine-tuned to mitigate performance degradation?

One of the challenges in setting up experiments to test and answer these questions is the lack of large datasets with accurate and extensive ground truth data on lighting and head pose. Thus a secondary set of research questions arises with regard to testing hypotheses on real data. More specifically, a real lighting dataset is introduced to address the following questions.

- 4) Does a synthetic dataset created by a relighting/pose-adapting GAN behave in the same way as a real multi-lighting/multi-pose dataset?
- 5) What are the reasons for this gap? What have we learned from these experiments?

As a first step toward answering these research questions, this work employs a SoA re-lighting methodology to augment a set of high-quality facial images with directional lighting effects. The effect of these augmentations on the performance of the ArcFace, which is a SoA FR model that ongoing provides the latest public code support, is quantified using

Receiver Operating Characteristic (ROC) curve techniques. A similar approach was used recently to validate synthetic facial identities [32]. Note that a re-lighting augmentation approach was adopted to evaluate the synthetic directional lighting data on the FR model. Then this work introduces a real public dataset, MultiPIE which includes more than 300 individual identities and 19 lighting scenes, to validate the results of this hypothesis on a real dataset with illumination variation. In this way, the robustness of the FR model to illumination is verified. Similarly, a SoA rotate-and-render technique is adopted to augment the face images with multiple head poses. The robustness of the FR model to head poses is verified by synthetic head poses and head poses from a real dataset with a limited identity number. This work also studies the feasibility of handling lighting variations by fine-tuning the neural FR network. The results for directional lighting are promising and indicate the potential for an end-to-end neural face authentication solution for in-the-wild faces.

The major contributions of this work are as follows:

- Evaluate the effect of directional facial lighting and multiple head poses on a state-of-the-art face recognizer.
- Fine-tuning alleviates the effect of direct light on the face recognizer.
- Public face datasets with illumination and head pose are quantified but concluded that they do not provide adequate variation in the context of extreme directional lighting conditions.

The optimization and compression of a CNN-based model for porting onto a neural accelerator require a complex and involved process of quantization and pruning of network layers, coupled with the extensive tuning of weights and validation experiments to retain model performance on one or more reference datasets [33], [34], [35], [36]. This work is not concerned explicitly with this process but rather seeks to establish if the original FR model can achieve sufficiently robust performance, thus justifying the effort involved in this model optimization and compression process.

The remainder of the paper is structured as follows. The literature survey is initially presented in Section II. Section III describes the datasets and methodology we adopted in our experiments. Subsequently, the initial experiments on GAN augmentations are described and discussed in Section IV and Section V. Further experiments and discussion about the inconsistent results between real and synthetic lighting are conducted in Section VI and Section VII. The fine-tuning strategies were proposed to compensate for the lighting variations in Section VIII. Finally, the findings of this research are summarized and discussed in Section IX.

## II. RELATED WORKS

### A. PORTRAIT RELIGHTING

Light Stage [37] captures the reflection field of faces by taking various photos of the same person under distinct lighting conditions, and subsequently re-rendering the 3D faces under novel lighting conditions. The method works

quite efficiently, but it requires expensive equipment and constrained experimental conditions, which is not suitable for consumer-level usage. This has promoted the development of single image portrait relighting without hardware conditions and constrained environments. Early approaches [38], [39] controlled lighting mainly by modifying the parameters of the directional lighting model, and the advent of deep learning has promoted the capability of authentic lighting effects. Reference [40] utilize deep neural networks to decompose image composition for relighting, however, it can only obtain coarse facial geometry, albedo and relighting results. Sun et al. [41] and Wang et al. [42] employed environment maps to build the illumination model, yielding more challenging results.

Zhou et al. [43] relight a facial image by adopting the Spherical Harmonics (SH) lighting model and using synthetic data to train the deep neural network. This model is able to produce renderings for arbitrary light directions that exhibit realistic shadows and specular highlights, and is able to generalize across a wide variety of subjects without introducing artifacts. The advance in the relighting techniques gives the ability to augment ‘in the wild’ face datasets introducing any variation in the direction, illuminance and intensity of the lighting while generating high-quality samples without artifacts. In our work, we evaluated the effect of directional illumination, and this model provides a solution for the limited variation of illuminated face datasets.

## B. FACE ROTATION

The face rotation, also known as synthetic multi-view faces, generates a face with a specified pose according to the input face. Face frontalization, which has the ability to recognize faces in various pose conditions, has attracted more attention in previous research. The existing methods can be roughly classified into 3D fitting based method [44], [45], encoder-decoder architecture based deep learning method [46], [47] and generative adversarial networks (GAN) based method [48], [49]. Most existing technologies require paired face images, and the quality of synthetic faces depends heavily on the training set, which limits the practical application of these technologies. Rotate-and-Render [50] proposed a novel self-supervised framework that can synthesize multiple views of faces using only single-view image collections in the wild. The key point of the rotate-and-render method is that they adopt a strong self-supervision by rotating faces in a 3D space and re-rendering these faces to a 2D plane. One of the aims in this work is to quantify the effect of arbitrary pose data on state-of-art neural face recognition. The rotate-and-render technology provides enough photo-realistic head pose data variations for us to answer the previous research questions.

## C. FACE RECOGNITION

With the advances in computational resources and with a surge in access to very large datasets, deep learning architectures have been developed and pushed the state-of-the-art in the FR task achieving exceptional accuracy results [7]. For example, the emergence of network architectures such

as VGGNet [51], GoogLeNet [52], ResNet [53], MobileFaceNet [54], FR3DNet [55] and etc. has improved the performance of face recognition. Another direction that has contributed to the development of face recognition in recent years has focused on exploring new approaches for learning loss functions for metrics including DeepFace [4], FaceNet [3], CosFace [6], ArcFace [5], CurricularFace [56], MagFace [57] and Elasticface [58], etc. Most of the loss functions are designed according to the criteria of maximizing the inter-class distance and minimizing the intra-class distance. More information regarding the state-of-the-art of deep neural FR approaches as well as the entire pipeline of the FR and the methods used are given in [25] and [59].

Despite the improvements, the FR task remained challenging in several cases. Studies revealed that many factors can have a negative impact on FR performance, with the main factors being pose, illumination, and others. Specifically, regarding the challenge of recognizing faces in various poses, the approaches employed include frontalizing faces [48], [49], building a mapping of side faces to frontal faces [60], combining head pose estimation with face recognition [61], [62], or using pose invariant face recognition (PIFR) [4], [63], [64]. On face samples with lighting variation, techniques were proposed that complement both the traditional and deep learning FR methods reporting improved performance [7], [65]. The approaches include pre-processing the facial samples to normalize any variation, before feeding it to the face recognition algorithm. This paper explores the impact of lighting and pose conditions on the robustness of state-of-the-art facial recognition (FR) techniques, and provides some insights into the study of end-to-end face recognition systems for multi-task scenarios such as extreme lighting conditions and facial poses.

## D. FACE DATASETS

Databases have been introduced to facilitate the development of FR models with light, pose, expression, and other variations. The Yale B dataset [66] contains 28 subjects under 9 poses and 64 illumination conditions in grayscale. The AR face dataset [67] was introduced with 4000 corresponding to 126 people’s faces in 1998. Their images are characterized by frontal faces with different facial expressions, illumination conditions, and occlusions. The CAS-PEAL face dataset [68] contains 99594 images of 1040 individuals, and each subject has 5 expressions, 6 accessories, and 15 lighting directions. All the subjects are from China. The PF01 (Postech Faces ’01) [69] contains true-color face images of 103 people with 4 illumination variations, and all of the subjects in the database are Asians. More recently the UHDB31 [70] was published having 21 poses and 3 lighting variations. Despite existing face datasets typically considering various lighting and pose conditions, they have various drawbacks such as limited human subjects, image resolution, ethnographic variation, and color variations. All of these datasets are not designed for portrait re-illumination or synthetic head pose purposes, so they do not contain multiple illuminations and

head pose information and are not able to represent 'in the wild' conditions.

### III. DATASETS AND METHODOLOGY

#### A. DATASET

##### 1) CelebA-HQ

CelebA-HQ [71] is a high quality face dataset derived from the CelebA [72] dataset. It contains 30k images of  $1024 \times 1024$  faces with roughly 6000 identities. This dataset is acquired by cropping faces from the original dataset by face landmarks and enhancing the image quality using a GAN trained super-resolution model. A bilinear interpolation and frame filter were used to scale the images to  $1024 \times 1024$ . The dataset was created and initially used to train PGAN [71] and StyleGAN [73]. CelebA-HQ is high quality synthetic dataset with sufficient subjects and limited lighting variation, which is a valid dataset for face recognition. In Section IV-A, CelebA-HQ is utilized to produce multiple directional illuminated face data using a face relighting method.

##### 2) BIWI

The Biwi Kinect Head Pose Dataset [74] contains around 15.8k images of 20 people (6 females and 14 males, 4 people were recorded twice). Each frame is provided with a depth image, corresponding RGB image and annotation. The resolution of each sample is  $640 \times 480$ , and the head pose range is yaw  $\in (-75, +75)$ , pitch  $\in (-60, +60)$  and roll  $\in (-50, +50)$ . Ground truth is provided in the form of the 3D position and rotation angle of the head. The BIWI is a 'real-world' dataset, which was captured by a Kinect at about one meter distance. BIWI is used in this work to generate synthetic head poses, though it has a limited number of identities. Then the synthetic head pose and the real head pose could be evaluated on the FR model.

##### 3) MultiPIE

The CMU Multi-PIE face database [75] contains 337 subjects, captured under 15 view points and 19 illumination conditions in four recording sessions for a total of more than 750,000 images. This dataset was published in 2013, which has a single background and a small span of illumination. Although it has some limitations, it has 337 subjects and 19 illumination conditions, making it in use by the research community. The cropped images from the original MultiPIE used for the experiment are shown in Figure 1.

#### B. FACE RELIGHTING METHOD

The lighting variation is applied to the CelebA-HQ dataset [71] and the MultiPIE dataset [75] via the Deep single image Portrait Relighting (DPR) technique [43]. In this method, a CNN is trained to generate a relighted image based on a Spherical Harmonics (SH) description of a lighting source. The method achieves SoA results, and in particular avoids introducing artifacts to the relighted samples - a drawback of other re-lighting methods [40], [42] that were

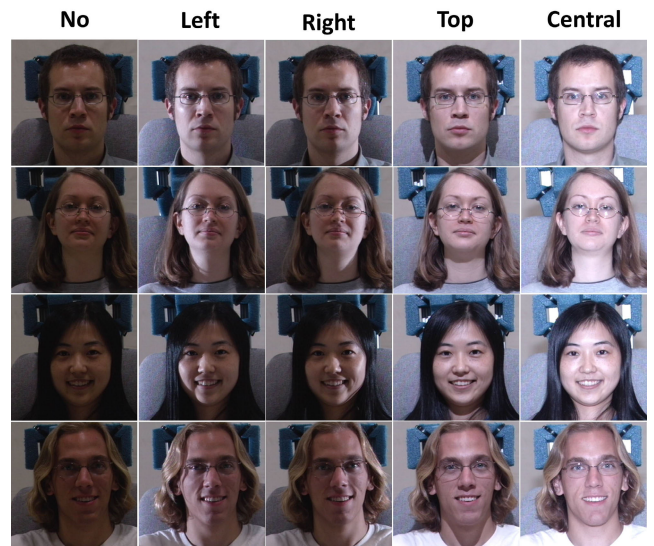


FIGURE 1. The original cropped images from CMU MultiPIE. The first column is for images without any flash (no lighting), and the rest of the column only used one flash (one lighting condition).

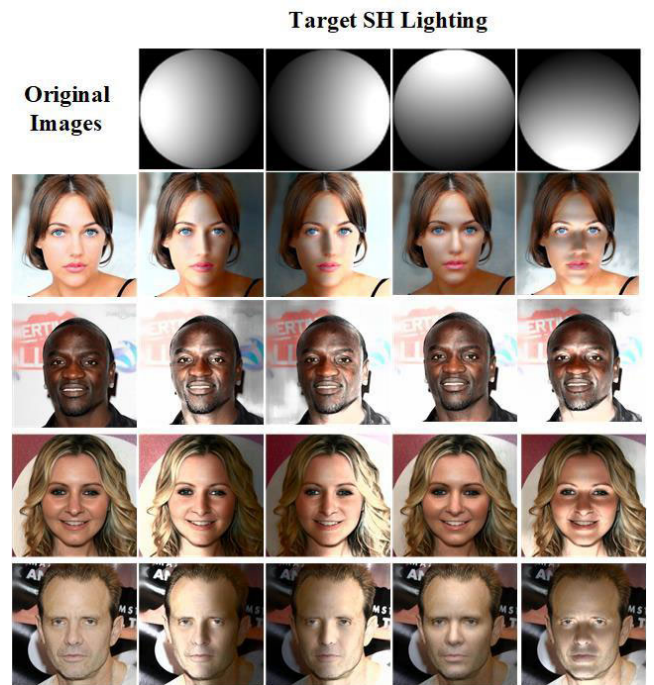


FIGURE 2. The target SH lighting that is examined in this work is presented on the top row. The original images of CelebA-HQ are on the left column. Examples of lighting injected in the original images using the DPR method [43] are shown for each examined illumination scenario (right, left, top & bottom).

considered for use in this study. The selected DPR method is trained on the well-known CelebA-HQ dataset which provides good variability in terms of subject identity, combined with consistent face image quality. This makes the combination of the DPR re-lighting methodology and CelebA-HQ ideal for this work as side effects are eliminated due to

either variable facial sample quality or re-lighting artifacts, either of which could distort our experimental outcomes. Note that DPR employs a generative adversarial network for re-lighting and obtains synthetic portraits with different directional lighting, so in this paper both synthetic lighting faces and GAN-generated lighting faces represent the portrait re-lighting.

In this work, experiments are restricted to a select set of directional lighting components in order to gain a better understanding of the overall effect of directional lighting. The selected scenarios that are examined include lighting from 4 main directions: right, left, top, and bottom of the face image. This has the added benefit of keeping the computation requirements for experiments bounded to a reasonable time frame, with most individual experiments completed in less than a 48-hour period.

The representative Spherical Harmonic (SH) lighting sources used are shown in Figure 2. More SH lighting scenarios can be found in [1]. The illumination variations (right, left, top and bottom) are introduced to each sample of the CelebA-HQ dataset, resolving with 4 new sets of the CelebA-HQ, each containing of one illumination variation (CelebA-HQ-right, CelebA-HQ-left, CelebA-HQ-top and CelebA-HQ-bottom). Examples of the CelebA-HQ samples after introducing the illumination variations are illustrated in Figure 2. It can be seen from Figure 2 that the DPR method has high quality outputs incorporating the target SH lighting to the images realistically and without generating any artifacts to the face images. Instructions on how to generate the sets of CelebA-HQ and MultiPIE with the different illumination scenarios are given in the Github repository of this work.<sup>2</sup>

### C. SYNTHETIC HEAD POSE METHOD

Rotate-and-render technique [50] is adopted to synthesize the head pose of the BIWI dataset. Specifically, 3D faces are generated from 2D using the 3D face fitting network 3DDFA [76] and a neural renderer. These faces are then rotated in 3D space and re-rendered to the 2D plane. An image-to-image translation network is then adjusted to fill in the invisible parts and get the output head pose. In this way, a single portrait image can be rotated to arbitrary angles, and the results eliminate artifacts caused by rotation, while existing local texture information can be preserved. Additionally, this technique is training face rotation in a fully unsupervised manner under in-the-wild scenarios. The above reasons make the rotate-and-render technique an ideal method for this work, which is strongly suitable in real-world circumstances.

The rotation of the head pose is divided into yaw, pitch and roll, as shown in Figure 3 [77]. Yaw is the rotation around the vertical axis, i.e., profile and frontal. Pitch is the rotation around the side-to-side axis, i.e., head up and head down and roll is planar rotation. Since the rotation of roll




FIGURE 3. Head rotation with Pitch, Yaw and Roll [77].

only involves the translation rotation of the image and the rotation of pitch faces in real-world are poor, we discuss the effect of yaw generated by the rotate-and-render technique in this research. We restrict our experiments to generating a series of fixed-angle faces to gain a better understanding of the synthetic head pose data. The synthesized pose also allows more granular control of pose angles than the real data, thus introducing an option for future research on head-pose to benefit from GAN-driven augmentation of frontal portrait images.


Examples of the CelebA-HQ samples after introducing the head pose variations are shown in Figure 4. It can be seen from Figure 4 that, subjectively, the rotation results have high quality, especially in rotation in a small range. In this work, the impact of face pose on the face recognition algorithm is further validated by the BIWI dataset, which is a real multi-view dataset obtained from a real environment. Since the BIWI dataset has limited identities and poses data, we use synthetic multi-view faces to further validate the impact of synthetic face pose on the face recognition algorithm. To get the best results when generating various head pose images, we select 'frontal' faces from the dataset as input to synthesize profile faces at arbitrary-angle faces. The synthesized pose also allows more granular control of pose angles than the real data - thus introducing an option for future research on head-pose to benefit from GAN-driven augmentation of frontal portrait images.

### D. ROC CURVES AS A METRIC

The Receiver Operating Characteristic (ROC) curve [78] illustrates the ability of a binary classifier system, which is created by plotting the true positive rate (TPR) against the false positive rate (FPR) at various threshold settings. X-axis, called false positive rate (FPR) or probability of false alarm, is defined as the proportion of negative data divided into positive categories. Y-axis, called true positive rate (TPR) or hit rate, means the proportion of positive data being divided into positive categories.

<sup>1</sup><https://zhoper.github.io/dpr.html>

<sup>2</sup><https://github.com/C3Imaging/Deep-Learning-Techniques/tree/Quantify-Retrain-FR-for-Light>



**FIGURE 4.** Examples of head pose variations generated by rotate-and-render method [50]. The images from the first column are original CelebA-HQ. The rest columns are the head pose with different angles.

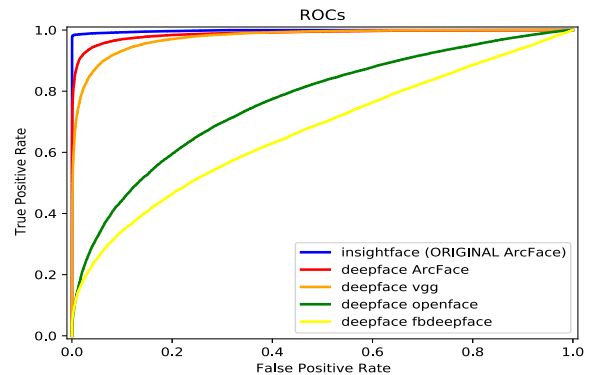
The ROC curve illustrates the performance of the classifier under different thresholds, which can visually demonstrate the uniqueness of the identity and is widely used to evaluate the performance of face recognition models [32]. Specifically, to compute a ROC curve, an equal number of positive-identity pairs (PP) and negative-identity pairs (NP) are created. Using the corresponding embeddings (extracted from the FR model for each image) of the pairs, similarity scores are calculated and used to plot the ROC curve.

The higher the TPR and the smaller the FPR, the better the classifier effect. Based on the performance of the ROC curve, the top-left curve of the ROC is close to 1 (i.e., the closer the TPR is to 1 when the FPR is close to 0), indicating that the performance of the FR model in the selected sample is more effective.

In this work, the ROC curve is used to evaluate the performance of the face recognition model. A PP is two images from the same identity and a NP is two images from distinct identities. The PPs are created by the whole dataset, and NPs are generated randomly with the same image pairs as PP. Then, the identity similarity score is calculated by the embeddings obtained from the face recognition model. True positive and true negative were obtained by calculating the identity similarity scores of PP. Similarly, false positive and false negative were obtained by using the identity similarity scores of NP. Then, TPR and FPR were calculated and the ROC curve is plotted.

**E. FACE RECOGNITION MODEL**

A public reference implementation of the ArcFace [5] model is available, as the authors have released optimized,



**FIGURE 5.** The ROCs comparison on different FR models.

pre-trained, weights for the model. This reference ArcFace model has high performance on the dataset used in this work and provides a useful public baseline for future performance comparisons. In general, any face recognition method could be used for experiments as long as they are state-of-the-art models and have a high performance on the datasets used in this work. Other SoA FR models such as FaceNet [3] or CosFace [6] do not provide reference implementations and thus restrict direct experimental comparisons. The unofficial, but public, implementation of Deepface framework<sup>3</sup> were also tested as shown in Figure 5 but could not provide a similar level of performance on the baseline or test datasets

<sup>3</sup><https://github.com/serengil/deepface>

used in this work. This has motivated our use of official ArcFace throughout this study.

ArcFace loss, or Additive Angular Margin Loss, is a loss function that maximizes the classification boundaries directly in the angular space used in face recognition tasks. In this work, the recommended workflow, by the authors of ArcFace is followed before the face samples are fed to the network. Firstly, the MTCNN [79] is used for detection. The detected area is cropped and resized to  $112 \times 112$ , using bilinear interpolation. Then the cropped faces are fed into the ArcFace network and 512-embedding will be computed corresponding to the faces. Finally, the identity similarity of two faces is obtained by calculating the cosine similarity, which is using two 512-embeddings. The pretrained network used in this work is provided by the authors of ArcFace and can be found in.<sup>4</sup>

Due to the introduction of lighting variation and pose variation, face detection is not able to process all the face images from the datasets. In the experiments, only the images which the face detection network was able to process in all the illumination scenarios and head pose conditions along with the original image, are used in order to keep the consistency in the experiments. A list of the images used in the experiments can be found in<sup>2</sup>.

#### IV. INITIAL EXPERIMENTS ON LIGHTING AND HEAD POSE

In this section, the experiments were conducted to quantify the effects of synthetic and real illumination conditions on FR performance, followed by experiments conducted to quantify the effects of synthetic and real pose conditions on FR performance.

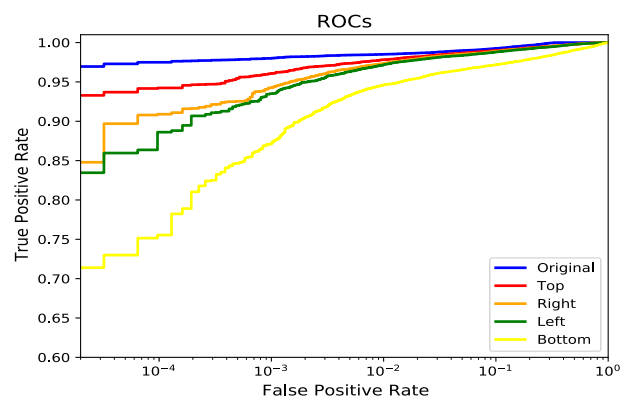
##### A. EXPERIMENTS WITH SYNTHETIC LIGHTING VARIATIONS

In this experiment, we evaluate the synthetic lighting variations from CelebA-HQ on the face recognition model. CelebA-HQ, as an initial base dataset, is utilized to produce multiple directional illuminated face data by the face re-lighting method.

The effects of the 4 directional lighting scenarios shown in Figure 2 on the FR's performance are examined. Initially, a ROC curve is calculated using only the samples from the test set of the original CelebA-HQ (ROC-Original) following the procedure described in section III-D. All possible PPs from the test set of the original CelebA-HQ are used, in total 31k image pairs and an equal number of NPs are created randomly. Using the corresponding FR embeddings, the similarity scores are calculated and used to plot the ROC curve.

The ROCs corresponding to re-lighting augmented scenarios is calculated following a similar procedure. The same positive and negative identity pairs as in ROC-Original are used but one of the samples from each pair has a re-lighting

<sup>4</sup><https://www.dropbox.com/s/ou8v3c307vyzawc/model-r50-arcface-ms1m-refine-v1.zip?dl=0>




**FIGURE 6.** ROC curves, representing the performance of the FR model<sup>4</sup> on the original images and the 4 initial directional illumination scenarios (left, right, top, bottom) examined in this work.

augmentation applied. This results in 4 main ROC curves (ROC-Left, ROC-Right, ROC-Top, ROC-Bottom) representing the FR's performance in each illumination scenario. The positive and negative pairs used to compute each ROC can be found in<sup>2</sup>. The resulting ROCs enable a direct comparison of the effects of different types of directional illumination with the original set of test image pairs and between them. This is presented in Figure 6.


From Figure 6 the initial experimental results are largely self-consistent and show well-defined performance degradation of the FR which is largely consistent with what might be expected. The ROC-Original curve illustrates that the FR model has a SoA performance on the non-augmented test dataset approaching close to unity, of 0.99 TPR on the corresponding to  $10^{-4}$  FPR value. The re-lighting augmented ROC curves show significant deviations from this baseline performance and are largely consistent with what might be expected. Thus, the smallest deviation is for the ROC-Top, which starts at 0.925 TPR, followed by the ROC-Right and ROC-Left curves at 0.86 and 0.85 respectively. The worst performing ROC is that of the bottom light, starting with a TPR of only 0.725.

Looking at the examples shown in Figure 2 these results make sense - the top lighting augmentation causes the least distortion to the facial image from a human perspective, whereas the bottom-lighting creates more obvious distortions in the facial features. Finally, the left/right lighting augmentations would be expected to have similar effects due to the symmetry of a human face. Note that the statistical variation between ROC-Left and ROC-Right is most likely due to statistical left-right pose variations in some facial samples leading to eccentricities in the corresponding lighting augmentations.

The initial results shown in Figure 6 encouraged a more extensive set of experiments to include additional top-right, top-left, and bottom-right, bottom-left lighting augmentations, to further improve our understanding of mixed directional lighting modes. The new SH lighting used and



**FIGURE 7.** The additional SH lighting that is examined in Figure 8 is presented on the top row. The original images of CelebA-HQ are in the left column. Examples of lighting injected in the original images using the DPR method [43] are shown for each examined illumination scenario (top-left, top-right, bottom-left, bottom-right).



**FIGURE 8.** Additional ROC curves, representing the performance of the FR model <sup>4</sup> on the original images and the 8 directional illumination scenarios (left, top-left, bottom-left, right, top-right, bottom-right, top and bottom) examined in this work.

examples of the CelebA-HQ samples after introducing these illumination variations are illustrated in Figure 7. The goals of this additional set of experiments were to provide a second validation of our results, in addition to exploring the effects of more varied re-lighting augmentations.

Due to the introduction of the new lighting variations, face detection is not able to process all the face images from the test set of the CelebA-HQ. Similarly as in section III-E, only the images which the face detection network was able to process in all 8 illumination scenarios and the original images

are used in order to keep the consistency in the experiments. Therefore, the initial test set of 8,654 images from 2k identities is reduced to 8,552 images from 1,979 identities. In order to calculate the ROCs, corresponding to the original images and the 8 illuminations scenarios the procedure described in III-D and IV-A, is followed, using the new tests. As the size of the test set is reduced, so is the number of all possible positive image pairs used to compute the ROCs. For this set of experiments 30k positive pairs and an equal number of negative pairs are used. These pairs can be found in <sup>2</sup>.

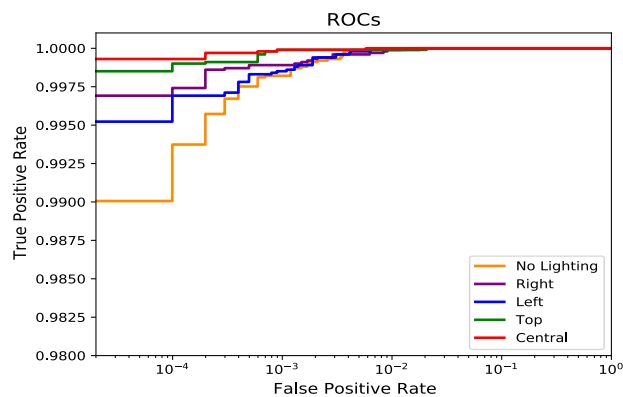
The primary directional ROCs curves (top, bottom, left, and right) presented in Figure 8, differ slightly from those of Figure 6, as the image pairs used in these experiments are different. However, their behavior has broadly similar characteristics. The bottom, bottom-left and bottom-right lighting augmentations are seen to be the most challenging for the FR task, while the top-left and top-right illuminations have the least effect on the FR’s performance. There is a small drop and a small increase in the FR’s performance from the top and right light respectively, compared to Figure 6. These are attributed to the use of different pairs. These results are useful as a baseline for the section VIII, as they help demonstrate that fine-tuning the primary set of directional lighting augmentations can generalize across a broader range of directional lighting effects.

**B. EXPERIMENTS WITH REAL LIGHTING VARIATIONS**

In this experiment, we evaluate the real lighting variations from MultiPIE on the face recognition model. Since MultiPIE has multiple views which may have an uncertain effect on the results, we only choose frontal faces as an initial experiment. Similar to the previous experiments, we create PPs and NPs at the beginning. First, two different frontal images with the same identity under ambient illumination conditions are employed to acquire a positive pair. Two frontal images with different identities under ambient illumination conditions are employed to acquire a negative pair. We can obtain 10059 positive pairs and subsequently generate the same number of negative pairs randomly from ambient illumination conditions. Then, for the other lighting conditions, we simply replace one image from the pair with a corresponding image having the relevant lighting condition. By applying the method in section III-D, we obtain the ROC curves of the real MultiPIE dataset in Figure 9.

The effects of 5 lighting conditions (ambient, central, left, right, and top lighting) shown in Figure 9, the initial experimental results show good performance across different lighting conditions. The ROC curves show that the FR model has a high performance on all conditions, which starts at 0.99 TPR. The ROC curve of ambient lighting condition is a bit poor. And the ROC curve of central lighting condition has good performance. This indicates that the ambient lighting condition is darker than others. These results appear to contradict the GAN-augmented results derived from CelebA-HQ, showing that the face recognition model can identify the original MultiPIE dataset very well. This is inconsistent with





**FIGURE 9.** ROC curves, representing the performance of the FR model <sup>4</sup> on the original MultiPIE images with the 4 illumination scenarios (the flashes we chose from MultiPIE are 00, 01, 07, 13, 16) examined in this work.

the results of Section IV-A, and some possible reasons for these outcomes will be further discussed in Section V and VI.

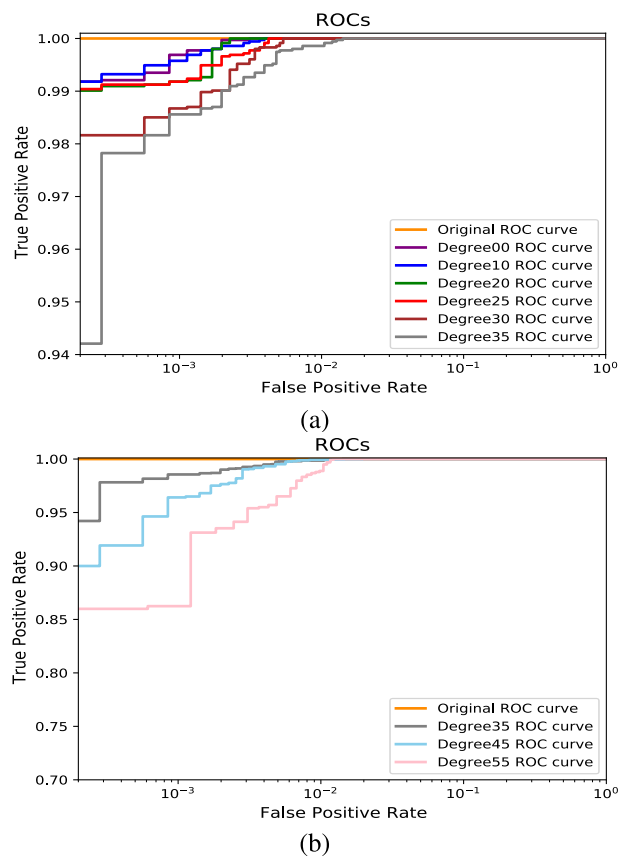
**C. EXPERIMENTS WITH SYNTHETIC HEAD POSE VARIATIONS**

This experiment quantifies the effect of synthetic multi-pose faces on face recognition. First, we validate the effect of yaw variation from 0 to 35 degrees on face recognition. Second, we extend the angular range of the BIWI dataset to study the effect of large yaw variation on face recognition.

The frontal faces are adopted to synthesize faces at 0, 10, 20, 25, 30, 35, 45, and 55 degrees with the rotate-and-render technique [50] to generate a large-scale synthetic head pose dataset. PPs with the same identity are generated using real frontal faces, and then NPs with the same number of different identities are generated to obtain 3.5K PPs and 3.5K NPs. The original ROC curve is plotted by using the PPs and NPs. For PP and NP, one image is from real frontal faces and the other image includes synthetic head pose variations. In this way, we can obtain the PPs and NPs of faces from the different head pose. Finally, the images from PPs and NPs are fed into ArcFace and the ROC curves are plotted via the method in Section III-D.

The resulting eight ROC curves (Degree00-ROC, Degree10-ROC, Degree20-ROC, Degree25-ROC, Degree30-ROC, Degree35-ROC, Degree45-ROC, and Degree55-ROC), respectively, represent the performance of face recognition under different pose variations. The plotted ROC curves enable a direct comparison of the effect of multi-pose faces with the original frontal image pairs, and between them. Figure 10 (a) shows the ROC curves drawn from 0 to 35 degree face pose, and (b) shows the ROC curves drawn from 35 to 55 degree face pose.

The BIWI original ROC curve in Figure 10 (a) shows that the face recognition model performs close to 1 on the original frontal dataset. The performance of Degree00 ROC curve and Degree10 ROC curve in Figure 10 (a) is almost unchanged and close to the performance of the original 'frontal' dataset



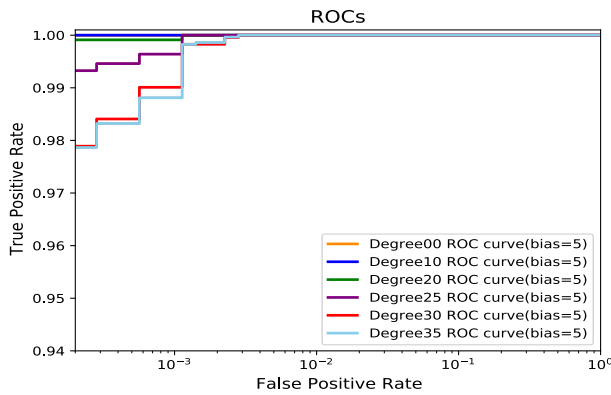
**FIGURE 10.** The ROC curves, representing the performance of the FR model <sup>4</sup> on the original images and the 8 synthetic head pose conditions examined in this work. Here the data are generated by the original BIWI, which yaw, pitch and roll  $\in (-5, 5)$  degrees.

on the FR model. This illustrates 10 degree head faces have almost no effect on FR performance. Then Degree 30 ROC curve notices a slight decrease in Figure 10 (a), the start of the ROC curve is still greater than 0.98 TPR.

Observing Figure 10 (b), the ROC curve at 35 degree, we can see that the accuracy starts to drop obviously. As the angle increases, the accuracy decreases more and more obvious. The start of the 45-degree ROC curve is 0.9 TPR and the 55-degree ROC curve is 0.86 TPR. It shows that the pose of 55 degrees is the most challenging one for the face recognition task in this experiment. This is consistent with our expectation that the face recognition performance decreases significantly with increasing input head pose, especially when the angle is greater than 35 degrees.

**D. EXPERIMENTS WITH REAL HEAD POSE VARIATIONS**

This experiment explores the use of ROC to measure and quantify the impact of real multi-pose head data on face recognition. While there is a wide range of pose variations in the BIWI dataset, this experiment explores the effect of yaw variations. We first limit the pitch and roll to a certain range to reduce their effect on the yaw. There is also a requirement that there are enough face images in the experiment to form



**FIGURE 11.** ROC curves, representing the performance of the FR model <sup>4</sup> on the original BIWI images with the 6 head pose examined in this work.

PPs and NPs. Restricted by these two factors, we define a certain range of varying angles as a fixed angle to increase the number of images, and this range is also called bias, which we set to 5 degrees. Also, we assume that faces have symmetry, and symmetrical faces have the same effect on face recognition. The pitch  $\in (-5, 5)$  and roll  $\in (-5, 5)$  remain constant during the experiment. As shown in Table 1, when yaw is greater than 35 degrees, the number and identities of images are less, and the effect of individual images on face recognition will have a significant impact on the ROC, causing the results to be unstable not to have reference value, so we only selected 0, 10, 20, 25, 30 and 35 degrees pose images in this experiment.

The method for generating PPs and NPs used to calculate face similarity is as follows. The PPs and NPs of the frontal faces are the same as Section IV-C, and approximately 3.5K PPs and 3.5K NPs could be obtained from the frontal faces. To maintain consistency in the experiment, the PPs for the remaining angles of the head pose consisted of one frontal face and one head pose (10, 20, 25, 30, and 35), respectively. In this way, 10 degree head poses have 4K PPs, 20 degree head poses can get 2.2K PPs, 25 degree head poses can generate 2.2K PPs, 30 degree head poses can get 2.3K PPs and 35 degree head poses can get 2.8K PPs. The image pairs from PPs and NPs were then fed together into ArcFace to calculate the similarity and get the TPR and FPR. Thus the ROC curves of real BIWI were plotted, as shown in Figure 11.

Figure 11 shows that the change in real-world pose will lead to a slight change in face recognition performance as expected. From Figure 11, we can see that frontal faces and 10-degree faces perform almost the same on the FR model, indicating that real data frontal and 10-degree head pose have almost no effect on FR performance. Observing the ROC curve at 20 and 25 degrees in Figure 11, we can find that the face recognition model can still achieve 99% accuracy. The ROC curve at 30 and 35 degrees in Figure 11 has an obvious drop, but the face recognition model can reach close to 98% accuracy. This phenomenon indicates that the face recognition model has high robustness for real faces

rotated within 35 degrees along the yaw direction. Compare to the synthetic head pose result in Figure 10 (a), it shows that the performance of the FR model on synthetic faces in Figure 10 (a) are consistent with the performance of FR model on real faces in Figure 11.

## V. DISCUSSION THE INITIAL EXPERIMENTS

Section IV illustrates and quantifies the effect of GAN-generated data and real data with lighting and pose variations on the face recognition model. First, the effect of synthetic illumination data and real illumination data on face recognition is evaluated, and the ROC curves show that the performance of synthetic illumination from CelebA-HQ and real illumination from MultiPIE are inconsistent on the SoA face recognition model. The performance of the ROC curve of real illumination is much better than that of synthetic illumination. Subsequently, the experimental results of Section IV-C and Section IV-D show that the performance of synthetic pose data and real pose data on the face recognition model is similar. The ROC curves illustrate the impact of various head pose data on the face recognizer and explore the behavior and capabilities of real pose samples compared to synthetic pose samples. Experiments have shown that the FR model is challenged by various pose conditions. Head pose ( $>35$  degrees) will cause degradation of the FR model, but the FR model is still robust enough even at 55 degrees for head poses. The synthetic pose dataset agrees with the real pose dataset, indicating that the synthetic pose dataset can replace or enhance the real dataset to some extent.

The relighting model is trained on CelebA-HQ, which has about 6000 identities. And the number of identities in MultiPIE is 337. The number of identities between MultiPIE and CelebA-HQ may lead to differences in ROC curves. Further work should focus on experiments to determine if the difference of scale between datasets affects the results, but lies outside of the scope of this work. The next set of experiments explores the potential probabilities for inconsistent results between real lighting and relighting portraits.

## VI. FURTHER EXPERIMENTS TO ILLUSTRATE THE RELIGHTING EFFECTS

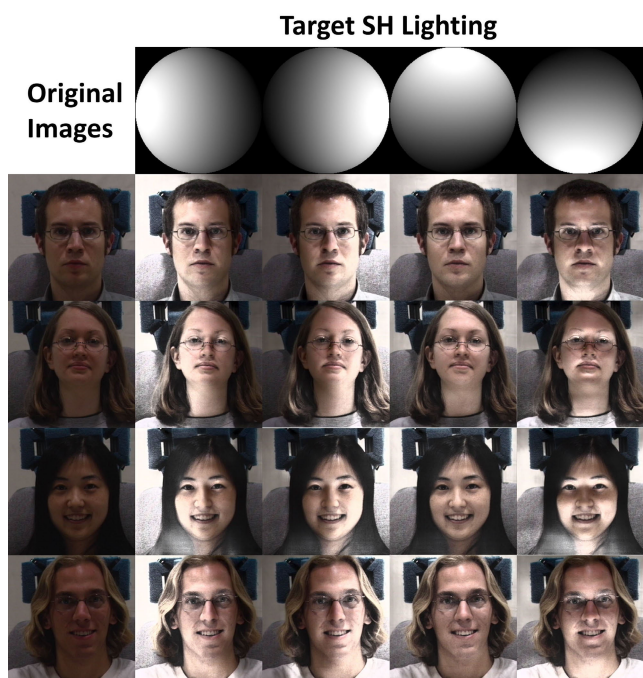
It is clear that GAN + CelebA (G+C) gives very different results from MultiPIE in Section IV-A and Section IV-B. The reasons for this inconsistent outcome are explored and clarified in this section. First, the effect of relighting portraits is quantified at MultiPIE on the face recognizer. Second, the discrepancy between samples of relighting faces and real-world faces is measured, as well as the potential explanations for that discrepancy.

### A. EXPERIMENTS WITH RELIGHTING VARIATIONS ON MultiPIE

In this section, the effect of relighting on MultiPIE performance is evaluated as a way to compare the effect of real lighting with synthetic lighting on face recognition performance. Firstly, ambient lighting images from MultiPIE are selected

**TABLE 1.** Define yaw angle of real BIWI (pitch  $\in (-5, +5)$  and roll  $\in (-5, +5)$ ).


Define Angle	0°	10°	20°	25°	30°	35°	45°	55°
Real Angle Range	$\pm 0^\circ \sim \pm 5^\circ$	$\pm 5^\circ \sim \pm 15^\circ$	$\pm 15^\circ \sim \pm 25^\circ$	$\pm 20^\circ \sim \pm 30^\circ$	$\pm 25^\circ \sim \pm 35^\circ$	$\pm 30^\circ \sim \pm 40^\circ$	$\pm 40^\circ \sim \pm 50^\circ$	$\pm 50^\circ \sim \pm 60^\circ$
Number of Images	293	300	177	148	152	164	143	36
Number of ID	20	19	13	13	13	11	9	5
Pairs	3539	4059	2236	2222	2323	2857	1755	693



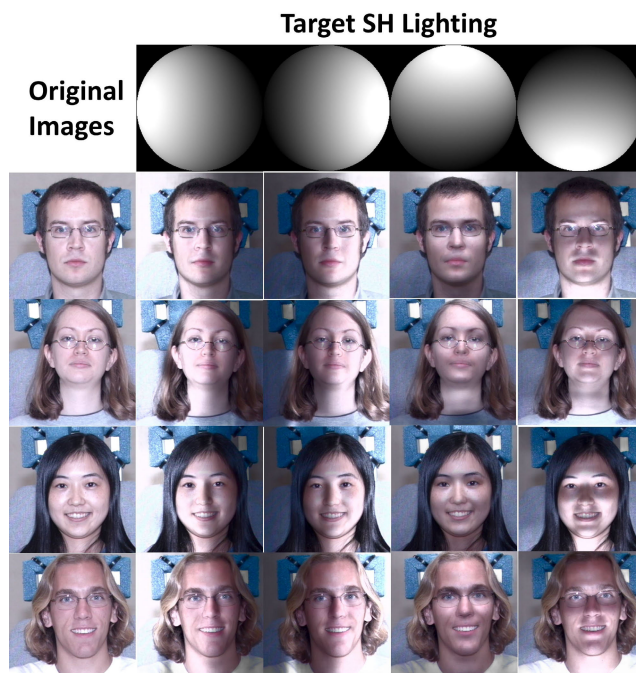
**FIGURE 12.** The target SH lighting that is examined in this work is presented on the top row. The original images of CMU MultiPIE from no lighting condition are shown in the left column. Examples of lighting injected in the original images using the DPR method [43] are shown for each examined illumination scenario (right, left, top and bottom).

as original images, and the samples of the relighting images are shown in Figure 12. Similar to the section IV-B, we create positive pairs and negative pairs from original ambient lighting images in MultiPIE. Subsequently, we simply replace one image from the pair with the relighting images. For example, a positive pair for directional right lighting is one image with ambient lighting from the original MultiPIE and another image with right lighting generated by relighting method. In this way, we could get one real pair and four relighting pairs. The relevant images are fed into the face recognition model and through calculation the directional lighting ROC curves from MultiPIE are plotted in Figure 13.

The ROCs in Figure 13 present the performance of the FR model on the original ambient lighting images from multiPIE and 4 directional lighting images. From Figure 13, the ROC-Top obtains almost the same performance as the ROC-Original on face recognition. Followed by a slight drop in the ROC-Right and ROC-Left, and the worst degradation is ROC-bottom, where the starting TPR is only slightly above 0.7. Compared to Figure 6, it is illustrated that the ROCs




**FIGURE 13.** ROC curves, representing the performance of the FR model <sup>4</sup> on the original MultiPIE images with no lighting and the 4 initial directional illumination scenarios (left, right, top and bottom) examined in this work.



**FIGURE 14.** The target SH lighting that is examined in this work is presented on the top row. The original images of CMU MultiPIE from central lighting condition are shown in the left column. Examples of lighting injected in the original images using the DPR method [43] are shown for each examined illumination scenario (right, left, top and bottom).

corresponding to the performance of the face recognition has a similar result with the initial ROCs from CelebA-HQ. Compared to Figure 9, the ROCs from the original MultiPIE are significantly better than the ROCs with relighting



**FIGURE 15.** ROC curves, representing the performance of the FR model <sup>4</sup> on the original MultiPIE images with central lighting and the 4 initial directional illumination scenarios (left, right, top and bottom) examined in this work.

method. From Figure 12, the relighting samples are dark and have many artifacts, far inferior to the synthetic results of CelebA-HQ in Figure 2.

We suspect that the initial ambient relighting images are too dark, which affects the outcome of the re-lighting portraits. Thus the central lighting images from MultiPIE are selected as initial images to generate the relighting samples, as shown in Figure 14. From Figure 14, the relighting portraits are brighter than the samples in Figure 9, and the artifacts disappear apparently. This indicates that the lighting conditions of initial images will affect the relighting results. The PPs and NPs from the original central lighting images are also created to calculate the ROC curve. Same with the last experiment, one of the images from PP and NP is replaced by relighting images, which are used to plot the ROC curves. Figure 15 illustrates the performance of the FR model on the original central lighting images from MultiPIE and 4 directional lighting images. The ROCs in Figure 15 show that the bottom lighting results are better than the one in Figure 13, but it is still inconsistent with the ROCs from the real lighting experiment.

### B. COMPARE THE SAMPLES FROM RELIGHTING METHOD WITH THE SAMPLES FROM REAL LIGHTING ON MultiPIE

To further elaborate on the differences between synthetic samples, we employ image quality evaluation and color histograms to detail the differences between synthetic samples and real samples.

#### 1) COMPARE WITH IMAGE QUALITY ASSESSMENT

To evaluate the effectiveness of GAN-generated re-lit images with real-world images, we compute the structural similarity index (SSIM) and peak signal-to-noise ratio (PSNR) between real-world images and relighting images. SSIM is a perception-based model that treats image degradation as a change in perceptible structural information and also incorporates important perceptual phenomena, including luminance and contrast. PSNR compares the absolute error between the

maximum signal and the background noise. Here we take real lighting images as the reference images and compute the SSIM and PSNR via reference images and the target relighting images.

From Figure 16, the quality of the re-lit images is degraded to varying degrees compared to the original images. The relighting images which are generated by original no lighting images have a better quality than the relit images which are generated by original central lighting images. From the second row of Figure 16, we could find the images generated by the original no-lighting images have more artifacts than the real images (first row). From the third row of Figure 16 (b), we could find that the relighting images generated by central lighting images have more realistic lighting compared with the second row of (b), but the human face is blurred and smoothed clearly in the third row of (b). This might be the reason why they have lower SSIM and PSNR than the second row.


#### 2) COMPARE WITH COLOR HISTOGRAM

The color histogram is a representation of the distribution of colors in an image. For digital images, the color histogram denotes the number of each pixel in a fixed list of color ranges that span the color space of the image, i.e., the set of all possible colors. The color histogram is a statistic that can be viewed as an approximation of a potentially continuous distribution of color values. This experiment measures the differences in the global distribution of colors between the real and synthetic images by comparing the differences in the color histograms.

Figure 17 shows the color histogram of two sample identities. For the histograms of real faces in the first column, the highest peaks are mainly concentrated between 0 and 100 levels, which means that the faces are relatively dark. The histograms of the real right lighting face and the histograms of the real left lighting face are consistent for the same face, indicating that the left and right illumination of the same face are symmetrical. The histograms of the top lighting face are uniform, implying that the brightness of the entire image is almost uniform.

From the second column in Figure 17, the histograms of GAN1 left lighting faces and GAN1 right lighting faces are similar. And the level of the highest peaks is 256, which means they are 'over-exposed' or 'saturated'. And the second highest peaks are closer to 0 levels than the histograms of the real images, implying that GAN1+lighting images are darker than the real lighting images. It indicates that there is significant ambient lighting in addition to the directional lighting in the images of the second column and the ambient illumination level is low from the input images of GAN1+lighting. This is consistent with the experimental conditions since the input images in the second column are images without illumination conditions.

From the third column in Figure 17, the level of the highest peaks is 256 from the histograms of left/right light which is similar to the second column. And some features in these

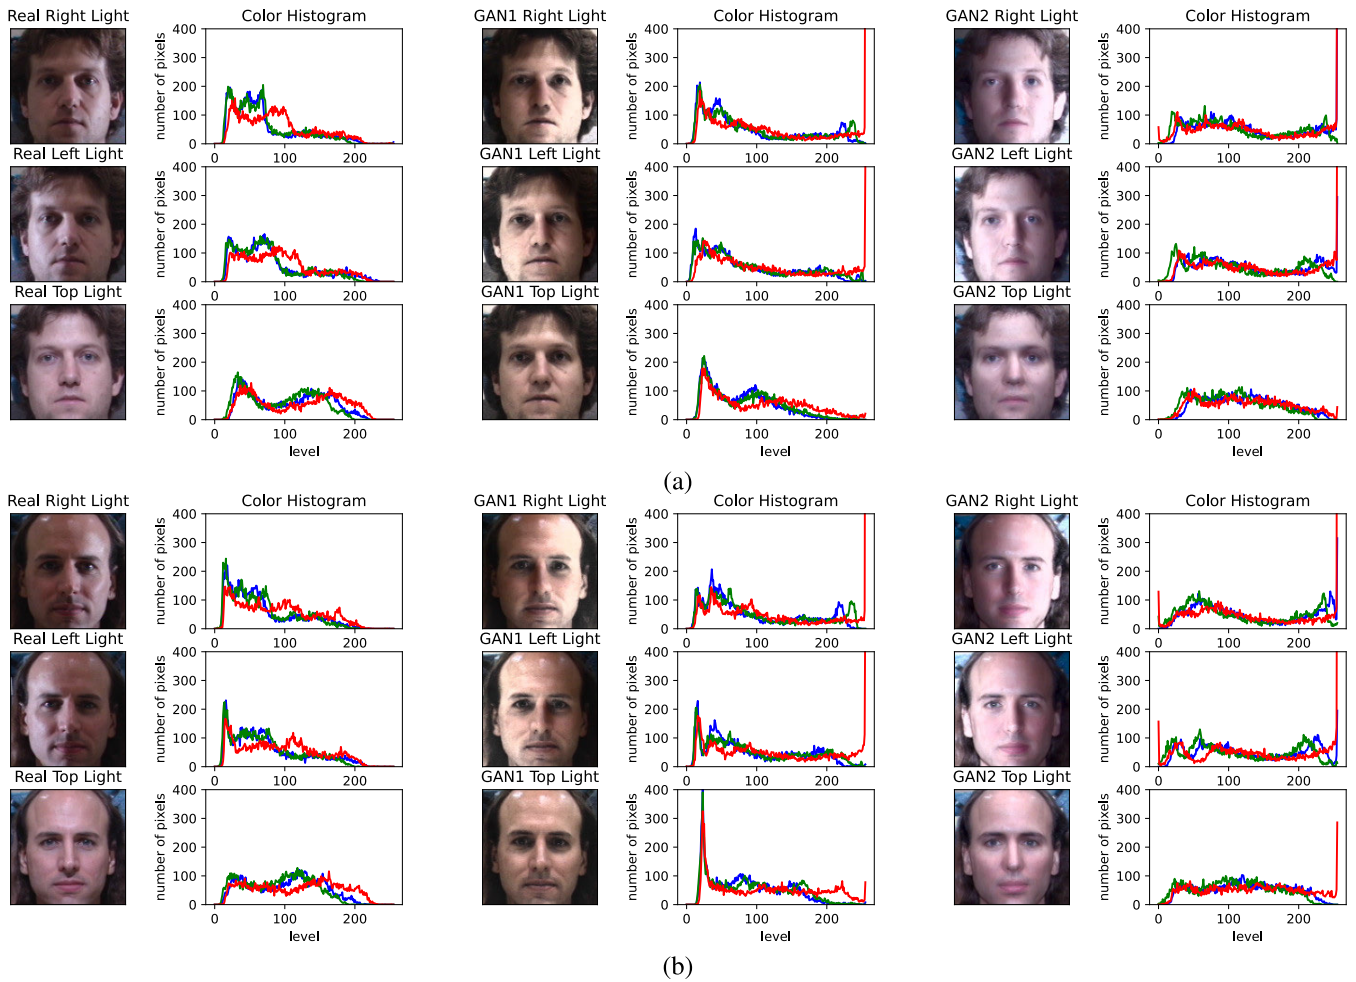


**FIGURE 16.** Two samples with different lighting conditions are presented here. The first row of each sample is the real lighting images from MultiPIE. The second row of each sample is the relighting images generated by the no-lighting images from the original MultiPIE. And the third row of each sample is the relighting images generated by the central images from the original MultiPIE. The first, third, and fifth columns are left lighting, right lighting, and top lighting conditions. The 2nd, 4th and 6th columns are their details respectively.

images such as eyes appear subjectively better than the ones in the second column because the initial image is a frontal illumination image. This comparison of real samples from the multi-PIE dataset with GAN-adjusted samples can only give us some subjective clues as to the origins of the observed

discrepancies in FR performance. A full resolution of this will need to be addressed in separate research work.

As we could see in Figure 18, the ambient light from the laboratory is at the top, and the walls of the laboratory may cause the reflections of flash lighting. This indicates that the



**FIGURE 17.** Two samples with color histograms are presented here. The first column of each sample is the cropped real lighting images from MultiPIE and their color histogram. The second column of each sample is the cropped relighting images generated by the ambient lighting images from the original cropped MultiPIE and their color histogram. And the third column of each sample is the cropped relighting images generated by the cropped central images from the original MultiPIE and their color histogram. The first, second, and third rows are right lighting, left lighting, and top lighting conditions.

mix of directional and ambient lighting is not controlled in Multi-PIE. It is hard to be sure how much of the directional lighting in the real data is ‘directional’ and how much is due to ambient reflections in the lab space. In addition, the GAN does cause some overexposure of the background pixels and the GAN images lack shadows - it is not clear how much overexposure is on the face or if this is enough to cause the observed deterioration in the ROC curve. These may be the reasons for the difference between the real images and relighting images.

**VII. DISCUSSION THE ADDITIONAL EXPERIMENTS**

Section VI illustrates that the effect of relighting faces from MultiPIE is consistent with the relighting faces from CelebA-HQ. So the difference between the effect of real lighting portraits and the effect of relighting portraits is not due to its scale or the underlying MultiPIE. We also gained an understanding that the input image of the relighting model will affect the re-illuminated portraits.



**FIGURE 18.** The collection laboratory of MultiPIE, and original images with ambient lighting.

From the comparison of the histograms, it shows that the color histogram is inconsistent between the real lighting faces and the relighting faces; it may imply that this is not only due to the GAN lighting model we used is not as good as

the real lighting, but also due to different mix of ambient and directional levels between the real lighting and GAN lighting.

In the first part of this work the potential effects of in-the-wild lighting conditions and multiple head pose, in particular directional lighting effects on a state-of-the-art neural face recognition method have been demonstrated and quantified. Thus we could get the conclusion that while GAN lighting has ‘issues’ the multi-PIE is not representative of true directional lighting either.

GAN lighting does give us extreme lighting conditions and it tells us the ‘worst’ effects we can expect. Next question is that should we be able to fine-tune for less significant distortions? As was indicated in [7] this problem is typically solved by additional pre-processing of image samples to correct for lighting conditions. Next experiment is to determine whether the FR can be fine-tuned to compensate for lighting effects.

### VIII. FINE-TUNING THE ArcFace

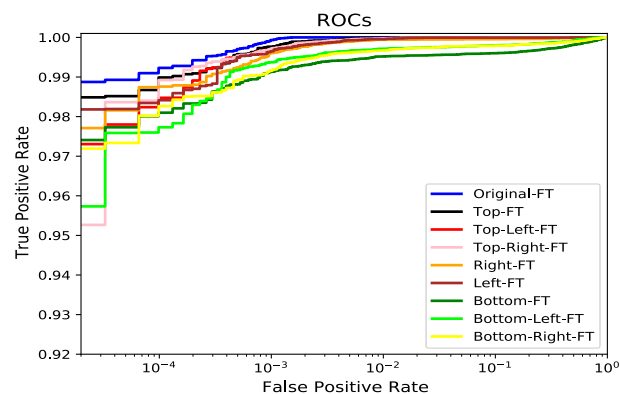
In this section the selected FR model is fine-tuned, using a similar approach to [80] with samples augmented with directional lighting.

#### A. FINE-TUNING PROCESS

The initial pretrained network provided by the authors of Arcface<sup>4</sup>, is fine-tuned using a training set comprising samples from the original CelebA-HQ dataset and samples with all 4 primary directional lighting augmentations (CelebA-HQ-Left, Right, Top and Bottom). In this experiment, 28,224 images were used from each set (CelebA-HQ-Left, Right, Top, Bottom and original) that the face detection network was able to process in all the illumination scenarios along with the original CelebA-HQ images. Then the dataset is divided into a train and test set with 19,570 images from 4k identities and 8,654 images from 2k identities, respectively. This is applied to each CelebA-HQ-set. Thus, 97,850 high-quality facial samples were used for fine-tuning, or 19,570 from the original data and each of the four primary lighting sub-categories.

For the re-training process, the standard Arcface loss function is used, with the learning rate set to 0.005 and a batch size of 128, following the instructions from the authors of ArcFace. The network is fine-tuned for 40 epochs, as the number of images used is relatively large and all network layers are unfrozen for the fine-tuning process. After 40 epochs the network showed satisfactory results on the training data used and therefore stopped. Longer training could result in over-fitting to the training data and thus not being able to generalize. More details regarding the fine-tuning process and the corresponding training code can be found at.<sup>5</sup> The fine-tuned network resulting from this re-training process is released at<sup>2</sup>.

The fine-tuned network is used to calculate the embeddings of the test samples. The same procedure as described in III-D and IV-A is followed to calculate the



**FIGURE 19. ROC curves, representing the performance of the finetuned FR network<sup>2</sup> on the original images and the 8 directional illumination scenarios (left, top-left, bottom-left, right, top-right, bottom-right, top and bottom) examined in this work.**

ROC-Original-finetuned (ROC-Original-FT) and the ROCs corresponding to the 8 different illumination scenarios (ROC-Left-FT, Right-FT, Top-FT, Bottom-FT, Top-Left-FT, Top-Right-FT, Bottom-Left-FT, Bottom-Right-FT,), using the positive and negative pairs from section IV-A. These ROCs are compared with the ROC-Original-FT and between them to explore whether the fine-tuned FR model is able to handle the variation in illumination as well as whether can generalize across the illuminations that were not used for the fine-tuning process.

#### B. ROCs COMPARISON

The ROCs representing the performance of the fine-tuned FR model on the original images and on the 8 directional lightings are presented in Figure 19. From Figure 19 it is illustrated that the ROCs corresponding to the fine-tuned FR model on the 4 main illuminations (left, right, top, bottom) used in the fine-tuning process (Figure 19) are at higher levels compared to the ROCs corresponding to the performance of the initial network (Figure 8) on these illumination scenarios. More importantly, the ROCs corresponding to the fine-tuned model on the 4 illuminations scenarios that are not used during the fine-tuning (top-left, top-right, bottom-left, bottom-right), are also at higher levels, thus showing that the network is able to generalize to other variations of illumination that it was not trained on. Overall, the performance of the fine-tuned FR model on any given illumination scenario has increased and its above 0.95 TPR, even on the lower FPR values. Notably, the ROCs are very close to the performance of the FR on the original images. Therefore, concluding that the FR model when trained with lighting variation is able to adapt and handle face samples that include illumination and achieve high accuracy results and also generalize across different illumination variations that are not used during fine-tuning. Thus, showing that the illumination can be compensated through training methods and augmentation techniques eliminating the need for pre-processing methods

<sup>5</sup><https://github.com/deepinsight/insightface>

to correct the lighting, which are not optimal for use in neural accelerators.

## IX. CONCLUSION AND FUTURE WORKS

It is clear from the results of the experiments, illustrated in Figure 10 and Figure 11 that a fully end-to-end neural face recognition model is robust to head pose. It is feasible to synthesize multiple head pose for large-scale training datasets in face recognition application. Through lighting experiments from Section IV and Section VI illustrated that there is still a gap between the current stage of the re-lighting method and the real environment. The potential reasons that lead to this difference may be the shadows in synthetic faces and a mixture of ambient and directional illumination in the original MultiPIE. The original images used for relighting experiments will affect the synthetic relighting results. In addition, SoA neural FR algorithms can be fine-tuned to handle difficult in-the-wild acquisition conditions such as directional lighting. In section VIII-B the practicality of fine-tuning a high-performing neural FR model has been demonstrated, recovering performance levels close to the original baseline for such lighting conditions. The fine-tuning process also indicated that generalization from the primary directions to combinations of directional lighting is achieved - a promising result given the non-linear nature of lighting conditions.

One of the reasons for the poor relighting results is shown in Figure 16 that the shadows in the synthetic images are inconsistent with the real images. Recent re-lighting method [81] deals with the hard shadows and relighting together, but only the nose shadow is considered and the experimental results we conducted on MultiPIE still have artifacts and blur, which affect the performance of the face recognition model. Other GAN relighting methods as we mentioned in the introduction, mostly concentrate on the environmental lighting condition, which is not in accordance with our work.

Future work we will improve the effect of shadows, and explore how shadows will affect the robustness of the face recognition model. We will collect face images with directional lighting conditions in the lab and real-world to produce a portrait lighting dataset. We also plan to generate a large synthetic dataset with various head pose and lighting conditions to explore the effect of combining pose and lighting on FR models. In addition, there are other challenges for FR algorithms in-the-wild, including those listed in the introduction. A broader study on factors that can affect FR is indicated. In this regard the availability of several large 3D facial model datasets [82] could provide sufficient individual identities and support more complex data variations to support such a study.

## REFERENCES

- [1] W. W. Bledsoe, "Some results on multicategory pattern recognition," *J. ACM*, vol. 13, no. 2, pp. 304–316, 1966.
- [2] S. Lawrence, C. L. Giles, A. C. Tsoi, and A. D. Back, "Face recognition: A convolutional neural-network approach," *IEEE Trans. Neural Netw.*, vol. 8, no. 1, pp. 98–113, Jan. 1997.
- [3] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 815–823.
- [4] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "DeepFace: Closing the gap to human-level performance in face verification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 1701–1708.
- [5] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 4690–4699.
- [6] H. Wang, Y. Wang, Z. Zhou, X. Ji, D. Gong, J. Zhou, Z. Li, and W. Liu, "CosFace: Large margin cosine loss for deep face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 5265–5274.
- [7] M. M. Ghazi and H. K. Ekenel, "A comprehensive analysis of deep learning based representation for face recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2016, pp. 34–41.
- [8] L. Tran, X. Yin, and X. Liu, "Disentangled representation learning GAN for pose-invariant face recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1415–1424.
- [9] J. Zhao, Y. Cheng, Y. Xu, L. Xiong, J. Li, F. Zhao, K. Jayashree, S. Pranata, S. Shen, J. Xing, S. Yan, and J. Feng, "Towards pose invariant face recognition in the wild," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 2207–2216.
- [10] S.-I. Choi, C.-H. Choi, and N. Kwak, "Face recognition based on 2D images under illumination and pose variations," *Pattern Recognit. Lett.*, vol. 32, no. 4, pp. 561–571, 2011.
- [11] A. Lanitis, C. J. Taylor, and T. F. Cootes, "Automatic interpretation and coding of face images using flexible models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 19, no. 7, pp. 743–756, Jul. 1997.
- [12] J. R. Beveridge, D. S. Bolme, B. A. Draper, G. H. Givens, Y. M. Lui, and P. J. Phillips, "Quantifying how lighting and focus affect face recognition performance," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. -Workshops*, Jun. 2010, pp. 74–81.
- [13] J.-Y. Zhu, W.-S. Zheng, F. Lu, and J.-H. Lai, "Illumination invariant single face image recognition under heterogeneous lighting condition," *Pattern Recognit.*, vol. 66, pp. 313–327, Jun. 2017.
- [14] J. W. Wang, N. T. Le, J. S. Lee, and C. C. Wang, "Illumination compensation for face recognition using adaptive singular value decomposition in the wavelet domain," *Inf. Sci.*, vol. 435, pp. 69–93, Apr. 2018.
- [15] A. Pena, A. Morales, I. Serna, J. Fierrez, and A. Lapedriza, "Facial expressions as a vulnerability in face recognition," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2021, pp. 2988–2992.
- [16] Y. Peng and H. Yin, "Facial expression analysis and expression-invariant face recognition by manifold-based synthesis," *Mach. Vis. Appl.*, vol. 29, no. 2, pp. 263–284, Feb. 2018.
- [17] M. Pavlovic, R. Petrovic, B. Stojanovic, and S. Stankovic, "Facial expression and lighting conditions influence on face recognition performance," in *Proc. ICETAN Conf.*, 2018, pp. 777–781.
- [18] S. Riaz, Z. Ali, U. Park, J. Choi, I. Masi, and P. Natarajan, "Age-invariant face recognition using gender specific 3D aging modeling," *Multimedia Tools Appl.*, vol. 78, no. 17, pp. 25163–25183, Sep. 2019.
- [19] D. Deb, L. Best-Rowden, and A. K. Jain, "Face recognition performance under aging," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jul. 2017, pp. 46–54.
- [20] C. Werther, M. Ferguson, K. Park, T. Kling, C. Chen, and Y. Wang, "Gender effect on face recognition for a large longitudinal database," in *Proc. IEEE Int. Workshop Inf. Forensics Secur. (WIFS)*, Dec. 2018, pp. 1–7.
- [21] N. Narang and T. Bourlai, "Gender and ethnicity classification using deep learning in heterogeneous face recognition," in *Proc. Int. Conf. Biometrics (ICB)*, Jun. 2016, pp. 1–8.
- [22] A. Nabatchian, E. Abdel-Raheem, and M. Ahmadi, "Illumination invariant feature extraction and mutual-information-based local matching for face recognition under illumination variation and occlusion," *Pattern Recognit.*, vol. 44, nos. 10–11, pp. 2576–2587, Oct. 2011.
- [23] L. Cimmino, M. Nappi, F. Narducci, and C. Pero, "M2FRED: Mobile masked face REcognition through periocular dynamics analysis," *IEEE Access*, vol. 10, pp. 94388–94402, 2022.
- [24] D. Zeng, R. Veldhuis, and L. Spreeuwiers, "A survey of face recognition techniques under occlusion," *IET Biometrics*, vol. 10, no. 6, pp. 581–606, Nov. 2021.
- [25] H. Du, H. Shi, D. Zeng, X.-P. Zhang, and T. Mei, "The elements of end-to-end deep face recognition: A survey of recent advances," *ACM Comput. Surveys*, vol. 54, no. 10s, pp. 1–42, Jan. 2022.




- [26] P. Corcoran, J. Lemley, C. Costache, and V. Varkarakis, "Deep learning for consumer devices and services 2—AI gets embedded at the edge," *IEEE Consum. Electron. Mag.*, vol. 8, no. 5, pp. 10–19, Sep. 2019.
- [27] A. Goel, C. Tung, Y.-H. Lu, and G. K. Thiruvathukal, "A survey of methods for low-power deep learning and computer vision," in *Proc. IEEE 6th World Forum Internet Things (WF-IoT)*, Jun. 2020, pp. 1–6.
- [28] F. Li, B. Liu, X. Wang, B. Zhang, and J. Yan, "Ternary weight networks," 2016, *arXiv:1605.04711*.
- [29] Z. He, B. Gong, and D. Fan, "Optimize deep convolutional neural network with ternarized weights and high accuracy," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Jan. 2019, pp. 913–921.
- [30] B. Fleischer, S. Shukla, M. Ziegler, J. Silberman, J. Oh, V. Srinivasan, J. Choi, S. Mueller, A. Agrawal, T. Babinsky, and N. Cao, "A scalable multi-TeraOPS deep learning processor core for AI training and inference," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2018, pp. 35–36.
- [31] W. Guicquero and A. Verdant, "Algorithmic enablers for compact neural network topology hardware design: Review and trends," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, Oct. 2020, pp. 1–5.
- [32] V. Varkarakis, S. Bazrafkan, G. Costache, and P. Corcoran, "Validating seed data samples for synthetic identities—methodology and uniqueness metrics," *IEEE Access*, vol. 8, pp. 152532–152550, 2020.
- [33] W. Zhang, N. Wang, L. Li, and T. Wei, "Joint compressing and partitioning of CNNs for fast edge-cloud collaborative intelligence for IoT," *J. Syst. Archit.*, vol. 125, Apr. 2022, Art. no. 102461.
- [34] S. Kalapothas, G. Flamis, and P. Kitsos, "Efficient edge-AI application deployment for FPGAs," *Information*, vol. 13, no. 6, p. 279, May 2022.
- [35] S. A. Bhalgaonkar, M. V. Munot, and A. D. Anuse, "Pruning for compression of visual pattern recognition networks: A survey from deep neural networks perspective," in *Pattern Recognition and Data Analysis With Applications*. Singapore: Springer, 2022, pp. 675–687.
- [36] U. Kulkarni, S. Meena, S. V. Gurlahosur, P. Benagi, A. Kashyap, A. Ansari, and V. Karnam, "AI model compression for edge devices using optimization techniques," in *Modern Approaches in Machine Learning and Cognitive Science: A Walkthrough: Latest Trends in AI*. Cham, Switzerland: Springer, 2021, pp. 227–240.
- [37] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar, "Acquiring the reflectance field of a human face," in *Proc. 27th Annu. Conf. Comput. Graph. Interact. Techn. (SIGGRAPH)*, 2000, pp. 145–156.
- [38] Y. Wang, Z. Liu, G. Hua, Z. Wen, Z. Zhang, and D. Samaras, "Face relighting from a single image under harsh lighting conditions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2007, pp. 1–8.
- [39] Y. Wang, L. Zhang, Z. Liu, G. Hua, Z. Wen, Z. Zhang, and D. Samaras, "Face relighting from a single image under arbitrary unknown lighting conditions," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 11, pp. 1968–1984, Oct. 2009.
- [40] S. Sengupta, A. Kanazawa, C. D. Castillo, and D. W. Jacobs, "SFS-Net: Learning shape, reflectance and illuminance of faces 'in the wild,'" in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 6296–6305.
- [41] T. Sun, J. T. Barron, Y.-T. Tsai, Z. Xu, X. Yu, G. Fyffe, C. Rhemann, J. Busch, P. E. Debevec, and R. Ramamoorthi, "Single image portrait relighting," *ACM Trans. Graph.*, vol. 38, no. 4, pp. 1–79, 2019.
- [42] Z. Wang, X. Yu, M. Lu, Q. Wang, C. Qian, and F. Xu, "Single image portrait relighting via explicit multiple reflectance channel modeling," *ACM Trans. Graph.*, vol. 39, no. 6, pp. 1–13, Dec. 2020.
- [43] H. Zhou, S. Hadap, K. Sunkavalli, and D. Jacobs, "Deep single-image portrait relighting," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 7194–7202.
- [44] Y. Feng, F. Wu, X. Shao, Y. Wang, and X. Zhou, "Joint 3D face reconstruction and dense alignment with position map regression network," in *Proc. ECCV*, 2018, pp. 534–551.
- [45] A. Jourabloo, M. Ye, X. Liu, and L. Ren, "Pose-invariant face alignment with a single CNN," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 3200–3209.
- [46] Y. Zhang, M. Shao, E. K. Wong, and Y. Fu, "Random faces guided sparse many-to-one encoder for pose-invariant face recognition," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 2416–2423.
- [47] M. Shao, Y. Zhang, and Y. Fu, "Collaborative random faces-guided encoders for pose-invariant face representation learning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 4, pp. 1019–1032, Apr. 2018.
- [48] R. Huang, S. Zhang, T. Li, and R. He, "Beyond face rotation: Global and local perception GAN for photorealistic and identity preserving frontal view synthesis," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2439–2448.
- [49] S. Zhang, Q. Miao, M. Huang, X. Zhu, Y. Chen, Z. Lei, and J. Wang, "Pose-weighted GAN for photorealistic face frontalization," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2019, pp. 2384–2388.
- [50] H. Zhou, J. Liu, Z. Liu, Y. Liu, and X. Wang, "Rotate-and-render: Unsupervised photorealistic face rotation from single-view images," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 5911–5920.
- [51] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014, *arXiv:1409.1556*.
- [52] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2015, pp. 1–9.
- [53] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [54] S. Chen, Y. Liu, X. Gao, and Z. Han, "MobileFaceNets: Efficient CNNs for accurate real-time face verification on mobile devices," in *Proc. Chin. Conf. Biometric Recognit.* Cham, Switzerland: Springer, 2018, pp. 428–438.
- [55] S. Z. Gilani and A. Mian, "Learning from millions of 3D scans for large-scale 3D face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 1896–1905.
- [56] Y. Huang, Y. Wang, Y. Tai, X. Liu, P. Shen, S. Li, J. Li, and F. Huang, "CurricularFace: Adaptive curriculum learning loss for deep face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 5901–5910.
- [57] Q. Meng, S. Zhao, Z. Huang, and F. Zhou, "MagFace: A universal representation for face recognition and quality assessment," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 14225–14234.
- [58] F. Boutros, N. Damer, F. Kirchbuchner, and A. Kuijper, "ElasticFace: Elastic margin loss for deep face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2022, pp. 1578–1587.
- [59] I. Adjabi, A. Ouahabi, A. Benzaoui, and A. Taleb-Ahmed, "Past, present, and future of face recognition: A review," *Electronics*, vol. 9, no. 8, p. 1188, Jul. 2020.
- [60] K. Cao, Y. Rong, C. Li, X. Tang, and C. C. Loy, "Pose-robust face recognition via deep residual equivariant mapping," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 5187–5196.
- [61] C. Bisogni, M. Nappi, C. Pero, and S. Ricciardi, "PIFS scheme for HEAd pose estimation aimed at faster face recognition," *IEEE Trans. Biometrics, Behav., Identity Sci.*, vol. 4, no. 2, pp. 173–184, Apr. 2022.
- [62] M. S. Sarfraz and O. Hellwich, "Head pose estimation in face recognition across pose scenarios," in *Proc. 3rd Int. Conf. Comput. Vis. Theory Appl.*, vol. 8, 2008, pp. 235–242.
- [63] C. Ding, J. Choi, D. Tao, and L. S. Davis, "Multi-directional multi-level dual-cross patterns for robust face recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 3, pp. 518–531, Mar. 2016.
- [64] K. Okada, S. Akamatsu, and C. Von der Malsburg, "Analysis and synthesis of pose variations of human faces by a linear PCMAP model and its application for pose-invariant face recognition system," in *Proc. 4th IEEE Int. Conf. Autom. Face Gesture Recognit.*, Sep. 2000, pp. 142–149.
- [65] H. Han, S. Shan, X. Chen, and W. Gao, "A comparative study on illumination preprocessing in face recognition," *Pattern Recognit.*, vol. 46, no. 6, pp. 1691–1699, Jun. 2013.
- [66] K.-C. Lee, J. Ho, and D. Kriegman, "Acquiring linear subspaces for face recognition under variable lighting," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 5, pp. 684–698, May 2005.
- [67] A. Martinez and R. Benavente, "The AR face database," CVC, Barcelona, Spain, Tech. Rep. 24, 1998.
- [68] W. Gao, B. Cao, S. Shan, X. Chen, D. Zhou, X. Zhang, and D. Zhao, "The CAS-PEAL large-scale Chinese face database and baseline evaluations," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 38, no. 1, pp. 149–161, Jan. 2008.
- [69] H. Dong and N. Gu. (2001). *Asian Face Image Database PF01*. [Online]. Available: <http://nava.postech.ac.kr/archives/irndb.html>
- [70] H. A. Le and I. A. Kakadiaris, "UHDB31: A dataset for better understanding face recognition across pose and illumination variation," in *Proc. IEEE Int. Conf. Comput. Vis. Workshops (ICCVW)*, Oct. 2017, pp. 2555–2563.
- [71] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of GANs for improved quality, stability, and variation," 2017, *arXiv:1710.10196*.

- [72] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 3730–3738.
- [73] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 4401–4410.
- [74] G. Fanelli, M. Dantone, J. Gall, A. Fossati, and L. Gool, "Random forests for real time 3D face analysis," *Int. J. Comput. Vis.*, vol. 101, no. 3, pp. 437–458, 2013.
- [75] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker, "Multi-PIE," *Image Vis. Comput.*, vol. 28, no. 5, pp. 807–813, 2010.
- [76] X. Zhu, X. Liu, Z. Lei, and S. Z. Li, "Face alignment in full pose range: A 3D total solution," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 1, pp. 78–92, Jan. 2019.
- [77] S. Basak, P. Corcoran, F. Khan, R. McDonnell, and M. Schukat, "Learning 3D head pose from synthetic data: A semi-supervised approach," *IEEE Access*, vol. 9, pp. 37557–37573, 2021.
- [78] A. Mansfield, *Information Technology—Biometric Performance Testing and Reporting—Part 1: Principles and Framework*, Standard ISO/IEC 19795-1, Apr. 1, 2006.
- [79] J. Xiang and G. Zhu, "Joint face detection and facial expression recognition with MTCNN," in *Proc. 4th Int. Conf. Inf. Sci. Control Eng. (ICISCE)*, Jul. 2017, pp. 424–427.
- [80] V. Varkarakis, S. Bazrafkan, and P. Corcoran, "Deep neural network and data augmentation methodology for off-axis iris segmentation in wearable headsets," *Neural Netw.*, vol. 121, pp. 101–121, Jan. 2020.
- [81] A. Hou, M. Sarkis, N. Bi, Y. Tong, and X. Liu, "Face relighting with geometrically consistent shadows," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 4217–4226.
- [82] H. Yang, H. Zhu, Y. Wang, M. Huang, Q. Shen, R. Yang, and X. Cao, "FaceScape: A large-scale high quality 3D face dataset and detailed rig-gable 3D face prediction," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 601–610.



**WANG YAO** (Graduate Student Member, IEEE) received the B.Sc. degree in computer science from Southwest University, China, in 2016, and the M.Sc. degree in control engineering from the University of Chinese Academy of Sciences (UCAS), in 2019. She is currently pursuing the Ph.D. degree with the University of Galway (NUIG). She is also an Intern at FotoNation/Xperi. Her research interest includes computer vision.



**VIKTOR VARKARAKIS** (Graduate Student Member, IEEE) received the B.Sc. degree in computer science and intelligent systems from the University of Piraeus, Greece, in 2017. He is currently pursuing the Ph.D. degree with the National University of Ireland Galway (NUIG). He is also with FotoNation/Xperi. His research interest includes machine learning using deep neural networks for tasks related to computer vision.



**GABRIEL COSTACHE** received the B.Sc. and M.Sc. degrees from the Faculty of Electronics and Telecommunications, Politehnica University of Bucharest, Romania, in 2003 and 2004, respectively, and the Ph.D. degree in image processing from the National University of Ireland Galway (NUIG), in 2006. He has been a part of FotoNation, since 2006. He is currently the Director of biometrics at FotoNation/Xperi, which develops technologies to process 2-D and 3-D imaging data.



**JOSEPH LEMLEY** (Member, IEEE) received the B.S. degree in computer science and the M.S. degree in computational science from Central Washington University, Ellensburg, WA, USA, in 2006 and 2016, respectively, and the Ph.D. degree from the National University of Ireland, Galway. He is currently Leading the Xperi's Sensing Group, which develops novel algorithms and artificial neural networks for upcoming sensor technologies. His research interests include artificial intelligence, deep learning, and computer vision. He received the 2017 Best Paper Joint Award for the *IEEE Consumer Electronics Magazine*, the Best Paper Second Place Award at ICCE 2018, and other awards in previous years.



**PETER CORCORAN** (Fellow, IEEE) holds the Personal Chair of electronic engineering at the College of Science and Engineering, University of Galway. He was a Co-Founder of several start-up companies, notably FotoNation, currently the Imaging Division of Xperi Corporation. He has over 600 technical publications and patents, over 100 peer-reviewed journal articles, 120 international conference papers, and a co-inventor of more than 300 granted U.S. patents. He is currently an IEEE fellow recognized for his contributions to digital camera technologies, notably in-camera red-eye correction and facial detection. He is a member of the IEEE Consumer Electronics Society for over 25 years. He is the Editor-in-Chief and the Founding Editor of *IEEE Consumer Electronics Magazine*.

• • •