

Towards End-to-End Neural Face Authentication in the Wild – Quantifying and Compensating for Directional Lighting Effects.

Viktor Varkarakis^a, Wang Yao^a, Peter Corcoran^a

^a*Department of Electronic Engineering, College of Engineering, National University of Ireland Galway, University Road, Galway, Ireland*

Abstract

The recent availability of low-power neural accelerator hardware, combined with improvements in end-to-end neural facial recognition algorithms provides enabling technology for on-device facial authentication. The present research work examines the effects of directional lighting on a State-of-Art (SoA) neural face recognizer. A synthetic re-lighting technique is used to augment data samples due to the lack of public data-sets with sufficient directional lighting variations. Top lighting and its variants (top-left, top-right) are found to have minimal effect on accuracy, while bottom-left or bottom-right directional lighting have the most pronounced effects. Following the fine-tuning of network weights, the face recognition model is shown to achieve close to the original Receiver Operating Characteristic curve (ROC) performance across all lighting conditions, and demonstrates an ability to generalize beyond the lighting augmentations used in the fine-tuning dataset. This work shows that a SoA neural face recognition models can be tuned to compensate for directional lighting effects, removing the need for a pre-processing step prior to applying facial recognition.

Keywords: Directional Lighting, Face Illumination, Face Recognition, Face Re-Lighting Method

1. Introduction

Human Face Recognition (FR) has been an active research field in computer vision since the early 1990's [38] and early Convolutional Neural Network (CNN) based approaches were in evidence before the end of that decade

[19]. 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 [32, 36, 7, 41, 26]. In much of the literature the test samples for FR are assumed to be 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-Art (SoA) FR systems.

The main factors that affect FR include (i) pose [37, 47, 3, 18], (ii) illumination [2, 49, 42], (iii) facial expression [29, 30, 28], (iv) age [31, 6] and, (v) gender variations [44, 27].

In this work, the focus is on the latest end-to-end fully neural FR techniques [9] as these represent current SoA in terms of accuracy and have the potential for implementation in the latest neural accelerators [4, 12]. The initial focus for implementation of neural algorithms in embedded devices was on network optimizations such as parameter quantization and pruning, compressed convolutional filters and matrix factorizations [12]. However, the attention has recently shifted towards specialized neural topologies [22, 16] and ultra-low power realizations in hardware [10, 14]. Such optimizations enable a SoA neural FR architectures to be implemented in a low-power consumer appliance, enabling a new generation of devices capable of identifying their owners, provide access control and personalize the device’s responses and behaviour. However, this introduces new challenges as such FR embodiment can no longer rely on pre-processing of input facial samples to optimize power consumption. Thus, all image processing must be achieved in a fully neural implementation, requiring a neural 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.

This leads to the research questions posed in this work: (i) can we better quantify the effects of the external factors that affect fully neural FR and develop metrics to evaluate these; and (ii) can a fully neural FR architecture be modified through tuning and/or re-training to compensate internally for such external factors? As human interaction with a consumer device will typically lead the user to look directly at a screen, or similar user-interface, this work has a focus on the effects of lighting/illumination scenarios on FR. A consideration of the effects of facial pose and other external factors is left

to future studies. Specifically for the lighting variation, numerous image pre-processing methods exist to improve the performance of the FR model [35, 50, 34], but studies exploring the tuning or training of FR models to be robust to lighting variations are relatively rare [5, 21].

As a first step towards 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 a SoA FR method is quantified using Receiver Operating Characteristic curve (ROC) techniques. A similar approach was used recently to validate synthetic facial identities [40]. Note that a re-lighting augmentation approach was adopted as existing public datasets do not provide sufficient lighting variability. This is discussed in the sections 4 and 6. Finally, this work 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.

2. Related Works

The advances in computational resources and with a surge in access to very large datasets, deep learning architectures have been developed and pushed the SoA in the FR task achieving exceptional accuracy results [26]. Famous deep neural approaches include DeepFace, FaceNet and ArcFace [32, 36, 7, 41]. Each work advanced the accuracy on the benchmarks of FR and new loss functions, pre-processing techniques and deep neural architectures were introduced. More information regarding the SoA of deep neural FR approaches as well as the entire pipeline of the FR and the methods used are given in [9, 1].

Despite the improvements, the FR task remained challenging in several cases. Studies revealed that many factors can have a negative impact on the FR performance, with the main factors being pose, illumination and others [2, 26, 18, 23]. Specifically, on face samples with lighting variation, techniques were proposed that compliment both the traditional and deep learning FR methods reporting improved performance [26, 15]. The approaches include pre-processing of the facial samples to normalise any variation, before feeding it to the face recognition algorithm [33, 17, 50, 34, 35].

Databases have been introduced to facilitate the development of FR models with light, pose, expression and other variations. The Yale b, CMU Multi-PIE, AR, Postech01 and UHDB31 [13, 25, 11, 8, 20] are a few of the datasets that have incorporate face variations and either focus on a single variation or a combination of different variations. Despite the developments of such datasets, the number of the variations and human subjects remain limited along with the fact that these datasets are usually acquired in controlled environments and therefore not being able to represent in-the-wild conditions.

Relighting techniques have been introduced in the literature with impressive results and able to introduce lighting into the face images without the degradation of the image by artifacts [48, 43] giving the ability to augment in-the-wild face datasets. Thus, providing a solution to the limited variation of lighting and human subjects in face datasets, which is discussed in section 4.

3. Methodology

In this section the various techniques and methodologies used in this work are detailed.

3.1. Face Re-Lighting Method

In this work, the lighting variation is applied to the CelebA-HQ dataset using the Deep Single Image Portrait Relighting (DPR) technique [48]. 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 that were considered for use in this study. The selected DPR method is trained on the well-known CelebA-HQ dataset which provides good variability in term 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.

In this work we restricted our experiments 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 reasonable time-frame, with most individual experiments completing in less than a 48 hour period.

The representative Spherical Harmonic (SH) lighting sources used are shown in Fig.1. 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 Fig.1. It can be seen from Fig.1 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 with the different illumination scenarios are given in the Github repository of this work ².

3.2. Face Recognition Model

A public reference implementation of the ArcFace [7] model is available, as the authors have released optimized, 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. This has motivated our use of ArcFace throughout this study. Other SoA FR models such as FaceNet [32] or CosFace [41] do not provide reference implementations and thus restrict direct experimental comparisons. An unofficial, but public, implementation of FaceNet ³ was also tested but could not provide a similar level of performance on the baseline or test datasets used in this work.

In this work, the recommended workflow, by the authors of ArcFace is followed before the face samples are fed to the network. The MTCNN [45] is used to detect and crop the face samples. The detected area is cropped and resized to 112×112 , using bilinear interpolation, before passing to the ArcFace network which calculates the 512-embedding corresponding to a facial sample. Finally, the cosine similarity is used to compute the score

¹<https://zhopper.github.io/dpr.html>

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

³<https://github.com/davidsandberg/facenet>




Figure 1: 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 [48] are shown for each examined illumination scenario (right, left, top & bottom).

representing the identity similarity when comparing two embeddings. The pretrained network used in this work is provided by the authors of ArcFace and can be found in ⁴.

Due to the introduction of the lighting variation and other factors, the face detection is not able to process all the face images from the CelebA-HQ sets. In the experiments only the images which the face detection network was able to process in all the illumination scenarios along with the original image, are used in order to keep the consistency in the experiments. Therefore, from the initial 30,000 images, in this work 28,222 are used from each set (CelebA-HQ-left, right, top, bottom and original). For the requirements of this study the dataset is divided into a train and test set with 19,570 images from 4k

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

identities and 8,654 images from 2k identities, respectively. This is applied to each CelebA-HQ-set. A list of the images used in the experiments can be found in ².

3.3. Using ROC Curves as a Metric

The ROC curve is created by plotting the true positive rate (TPR) against the false positive rate (FPR) at various threshold settings [24]. More specifically, to compute an ROC curve, an equal number of positive-identity and negative-identity pairs 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 closer an ROC curve is to unity, the better the performance of the FR model on the selected samples. More information regarding the ROC and its interpretation and use can be found in [40].

4. Quantifying the Effects of Illumination on the Face Recognition

In this section, the experiments conducted to quantify the effect of different lighting conditions on the FR’s performance are presented with a discussion of the findings.

4.1. Experiments on Initial Lighting Variations

The effects of the 4 directional lighting scenarios shown in Fig.1 on the FR’s performance are examined. Initially an 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 3.3. All possible positive pairs from the test set of the original CelebA-HQ are used, in total 31k image pairs and an equal number of negative pairs 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 are 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 augmentation applied. Thus resulting 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 ². 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 Fig.2.

From Fig.2 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 a TPR of only 0.725. Looking at the examples shown in Fig.1 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 slight variation between ROC-Left and ROC-Right is most likely due to slight left-right pose variations in some facial samples leading to eccentricities in the corresponding lighting augmentations.

4.2. Experiments on Additional Lighting Variations

The initial results shown in Fig.2 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 examples of the CelebA-HQ samples after introducing these illumination variations are illustrated in Fig.3. The goals of this additional set of experiments were to provide a second validation of our results, and in addition to explore the effects of more varied re-lighting augmentations.

Due to the introduction of the new lighting variations, the face detection is not able to process all the face images from the test set of the CelebA-HQ. Similarly as in section 3.2, 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




Figure 2: ROC curves, representing the performance of the FR model ⁴ on the original images and the 4 initial directional illumination scenarios (left, right, top, bottom) examined in this work.

original images and the 8 illuminations scenarios the procedure described in 3.3 and 4.1, 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 ².

The primary directional ROCs curves (top, bottom, left & right) presented in Fig.4, differ slightly from those of Fig.2, as the image pairs used in these experiments are different. However their behaviour has a broadly similar characteristic. 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 on the FR’s performance from the top and right light respectively, compared to the Fig.2. These are attributed to the use of different pairs. These results are useful as a baseline for the section 5, as they help demonstrate that fine-tuning on the primary set of directional lighting augmentations can generalize across a broader range of directional lighting effects.




Figure 3: The additional SH lighting that is examined in section 4.2 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 [48] are shown for each examined illumination scenario (top-left, top-right, bottom-left, bottom-right).

4.3. Experiments on Public Face Illumination Data-sets

Similar experiments as in 4.1 and 4.2, were conducted using face datasets which include illumination variation in their samples, to measure the effect of their lighting scenarios on the FR’s performance. More specifically the AR and Postech01 [25, 8] face datasets were used. These experiments did not show any measurable degradation of the FR’s performance across the illumination conditions provided in these data-sets. This is attributed to the relatively limited variation of the illumination scenarios, the numbers of human subjects, the total number of images and the controlled environments in which these face datasets were acquired. This is also indicated through the information regarding these datasets, provided in Table 1. Thus, revealing the need for the development of a face dataset that resembles in-the-wild conditions with illumination variation, which is further discussed in section 6.




Figure 4: Additional ROC curves, representing the performance of the FR model ⁴ 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.

Table 1: Summary Information of Face Data-sets with Illumination Variations. (The number of Pose Variations in this table is referred with regards to the Illumination Variation)

Face sets	Data-	#IDs	#Illumination Variations	#Pose Variations
AR [25]		126	3	1
Postech01 [8]		103	4	1

5. Fine-Tuning the Face Recognition model on Directional Illuminations

In the first part of this work the potential effects of in-the-wild lighting conditions, in particular directional lighting effects on a SoA neural face recognition method have been demonstrated and quantified. The next step is to determine whether the FR can be tuned to compensate for these effects. In this section the selected FR model is fine-tuned, using a similar approach to [39] with samples augmented with directional lighting.

5.1. Fine-tuning Process & ROCs Computation

The initial pretrained network provided by the authors of Arcface ⁴, 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 total 97,850 high-quality facial samples were used for fine-tuning, or 19,750 from the original data and each of the four primary lighting sub-category.

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 generalise. More details regarding the fine-tuning process and the corresponding training code can be found at ⁵. The fine-tuned network resulting from this re-training process is released at ².

The fine-tuned network, is used to calculate the embeddings of the test samples. The same procedure as described in 3.3 and 4.1 is followed to calculate the 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 4.2. 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 generalise across the illuminations that were not used for the fine-tuning process.

5.2. 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 the Fig.5. From Fig.5 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 (Fig.5) are at higher levels compared to the ROCs

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




Figure 5: ROC curves, representing the performance of the finetuned FR network ² 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.

corresponding to the performance of the initial network (Fig.4) 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 generalise 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 adopt and handle face samples that include illumination and achieve high accuracy results and also generalise 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 correcting the lighting, which are not optimal for use in neural accelerators.

6. Conclusion & Future Work

It is clear from the results of the experiments, illustrated in Fig. 2 & 4 that fully end-to-end neural FR solutions will be challenged by in-the-wild lighting conditions. As was indicated in [26] this problem is typically solved by additional pre-processing of image samples to correct for lighting conditions. In section 5.2 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. Note that providing a broader and more varied range of re-lighting samples and refining the training methodology to identify the more sensitive network layers in the ArcFace model should further improve these results, but even as they stand it is clear that a full end-to-end neural FR can be realized.

These initial results, especially the effectiveness of the fine-tuning process are very promising. They suggest that SoA neural FR algorithms can be fine-tuned to handle difficult in-the-wild acquisition conditions such as directional lighting. However 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 [46] could provide sufficient individual identities and support more complex data variations to support such a study.

Acknowledgments

This research is funded by (i) the Science Foundation Ireland Strategic Partnership Program (Project ID: 13/SPP/I2868), (ii) Irish Research Council Enterprise Partnership Ph.D. Scheme (Project ID: EPSPG/2020/40) and, (iii) Xperi Corporation, Ireland.

References

- [1] Adjabi, I., Ouahabi, A., Benzaoui, A., Taleb-Ahmed, A., 2020. Past, present, and future of face recognition: A review. *Electronics* 9, 1188.

- [2] Beveridge, J.R., Bolme, D.S., Draper, B.A., Givens, G.H., Lui, Y.M., Phillips, P.J., 2010. Quantifying how lighting and focus affect face recognition performance, *IEEE*. pp. 74–81.
- [3] Choi, I.S., Choi, C.H., Kwak, N., 2011. Face recognition based on 2d images under illumination and pose variations. *Pattern Recognition Letters* 32, 561–571. doi:10.1016/j.patrec.2010.11.021.
- [4] Corcoran, P., Lemley, J., Costache, C., Varkarakis, V., 2019. Deep learning for consumer devices and services 2—ai gets embedded at the edge. *IEEE Consumer Electronics Magazine* 8, 10–19. doi:10.1109/mce.2019.2923042.
- [5] Crispell, D., Biris, O., Crosswhite, N., Byrne, J., Mundy, J.L., 2017. Dataset augmentation for pose and lighting invariant face recognition. *arXiv preprint arXiv:1704.04326* .
- [6] Deb, D., Best-Rowden, L., Jain, A.K., 2017. Face recognition performance under aging, pp. 548–556. doi:10.1109/CVPRW.2017.82.
- [7] Deng, J., Guo, J., Xue, N., Zafeiriou, S., 2019. Arcface: Additive angular margin loss for deep face recognition, pp. 4690–4699.
- [8] Dong, H., Gu, N., 2001. Asian face image database pf01. 2006-01-05](2014-11). <http://nava.postech.ac.kr/archives/irndb.html> .
- [9] Du, H., Shi, H., Zeng, D., Mei, T., 2020. The elements of end-to-end deep face recognition: A survey of recent advances. *arXiv preprint arXiv:2009.13290* .
- [10] Fleischer, B., Shukla, S., Ziegler, M., Silberman, J., Oh, J., Srinivasan, V., Choi, J., Mueller, S., Agrawal, A., Babinsky, T., Cao, N., Chen, C.Y., Chuang, P., Fox, T., Gristede, G., Guillorn, M., Haynie, H., Klaiber, M., Lee, D., Lo, S.H., Maier, G., Scheuermann, M., Venkataramani, S., Vezyrtzis, C., Wang, N., Yee, F., Zhou, C., Lu, P.F., Curran, B., Chang, L., Gopalakrishnan, K., 2018. A scalable multi-teraops deep learning processor core for ai training and inference, pp. 35–36. doi:10.1109/VLSIC.2018.8502276.
- [11] Georgiades, A.S., Belhumeur, P.N., Kriegman, D.J., 2001. From few to many: Illumination cone models for face recognition under variable

lighting and pose. *IEEE transactions on pattern analysis and machine intelligence* 23, 643–660.

- [12] Goel, A., Tung, C., Lu, Y.H., Thiruvathukal, G.K., 2020. A survey of methods for low-power deep learning and computer vision, Institute of Electrical and Electronics Engineers Inc. doi:10.1109/WF-IoT48130.2020.9221198.
- [13] Gross, R., Matthews, I., Cohn, J., Kanade, T., Baker, S., 2010. Multiple. *Image and Vision Computing* 28, 807–813.
- [14] Guicquero, W., Verdant, A., 2020. Algorithmic enablers for compact neural network topology hardware design: Review and trends, pp. 1–5. doi:10.1109/iscas45731.2020.9181005.
- [15] Han, H., Shan, S., Chen, X., Gao, W., 2013. A comparative study on illumination preprocessing in face recognition. *Pattern Recognition* 46, 1691–1699.
- [16] He, Z., Gong, B., Fan, D., 2018. Optimize deep convolutional neural network with ternarized weights and high accuracy. *arXiv* .
- [17] Hussain Shah, J., Sharif, M., Raza, M., Murtaza, M., Ur-Rehman, S., 2015. Robust face recognition technique under varying illumination. *Journal of applied research and technology* 13, 97–105.
- [18] Lanitis, A., Taylor, C., Cootes, T., 1997. Automatic interpretation and coding of face images using flexible models. *Pattern analysis and machine intelligence, IEEE Transactions on* 19, 743–756. doi:10.1109/34.598231.
- [19] Lawrence, S., Giles, C.L., Tsoi, A.C., Back, A.D., 1997. Face recognition: A convolutional neural-network approach. *IEEE Transactions on Neural Networks* 8, 98–113. doi:10.1109/72.554195.
- [20] Le, H.A., Kakadiaris, I.A., 2017. Ufdb31: A dataset for better understanding face recognition across pose and illumination variation, pp. 2555–2563.
- [21] Le, H.A., Kakadiaris, I.A., 2019. Illumination-invariant face recognition with deep relit face images, in: 2019 IEEE Winter Conference on Applications of Computer Vision (WACV), IEEE. pp. 2146–2155.

- [22] Li, F., Zhang, B., Liu, B., 2016. Ternary weight networks .
- [23] Mahmood, Z., Ali, T., Khan, S.U., 2016. Effects of pose and image resolution on automatic face recognition. *IET biometrics* 5, 111–119.
- [24] Mansfield, A., 2006. Information technology–biometric performance testing and reporting–part 1: Principles and framework. *ISO/IEC* , 19791–19795.
- [25] Martinez, A.M., 1998. The ar face database. *CVC Technical Report*24 .
- [26] Mehdiipour Ghazi, M., Kemal Ekenel, H., 2016. A comprehensive analysis of deep learning based representation for face recognition, in: *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 34–41.
- [27] Narang, N., Bourlai, T., 2016. Gender and ethnicity classification using deep learning in heterogeneous face recognition. doi:10.1109/ICB.2016.7550082.
- [28] Pavlović, M., Petrović, R., Stojanović, B., Stanković, S., 2018. Facial expression and lighting conditions influence on face recognition performance. *Technical Report*.
- [29] Peña, A., Serna, I., Morales, A., Fierrez, J., Lapedriza, A., 2020. Facial expressions as a vulnerability in face recognition. *arXiv preprint arXiv:2011.08809* .
- [30] Peng, Y., Yin, H., 2018. Facial expression analysis and expression-invariant face recognition by manifold-based synthesis. *Machine Vision and Applications* 29, 263–284. doi:10.1007/s00138-017-0895-6.
- [31] Riaz, S., Ali, Z., Park, U., Choi, J., Masi, I., Natarajan, P., 2019. Age-invariant face recognition using gender specific 3d aging modeling. *Multimedia Tools and Applications* doi:10.1007/s11042-019-7694-1.
- [32] Schroff, F., Kalenichenko, D., Philbin, J., 2015. Facenet: A unified embedding for face recognition and clustering, pp. 815–823.
- [33] Shan, S., Gao, W., Cao, B., Zhao, D., 2003. Illumination normalization for robust face recognition against varying lighting conditions, *IEEE*. pp. 157–164.

- [34] Shinwari, A.R., Balooch, A.J., Alariki, A.A., Abdulhak, S.A., 2019. A comparative study of face recognition algorithms under facial expression and illumination, IEEE. pp. 390–394.
- [35] Ruiz-del Solar, J., Quinteros, J., 2008. Illumination compensation and normalization in eigenspace-based face recognition: A comparative study of different pre-processing approaches. *Pattern Recognition Letters* 29, 1966–1979.
- [36] Taigman, Y., Yang, M., Ranzato, M., Wolf, L., 2014. Deepface: Closing the gap to human-level performance in face verification, pp. 1701–1708.
- [37] Tran, L., Yin, X., Liu, X., 2017. Disentangled representation learning gan for pose-invariant face recognition, pp. 1283–1292. doi:10.1109/CVPR.2017.141.
- [38] Turk, M., Pentland, A., 1991. Eigenfaces for recognition. *Journal of cognitive neuroscience* doi:10.1162/jocn.1991.3.1.71.
- [39] Varkarakis, V., Bazrafkan, S., Corcoran, P., 2020a. Deep neural network and data augmentation methodology for off-axis iris segmentation in wearable headsets. *Neural Networks* 121, 101–121.
- [40] Varkarakis, V., Bazrafkan, S., Costache, G., Corcoran, P., 2020b. Validating seed data samples for synthetic identities—methodology and uniqueness metrics. *IEEE Access* 8, 152532–152550.
- [41] Wang, H., Wang, Y., Zhou, Z., Ji, X., Gong, D., Zhou, J., Li, Z., Liu, W., 2018a. Cosface: Large margin cosine loss for deep face recognition, pp. 5265–5274.
- [42] Wang, J.W., Le, N.T., Lee, J.S., Wang, C.C., 2018b. Illumination compensation for face recognition using adaptive singular value decomposition in the wavelet domain. *Information Sciences* 435, 69–93. doi:10.1016/j.ins.2017.12.057.
- [43] Wang, Z., Yu, X., Lu, M., Wang, Q., Qian, C., Xu, F., 2020. Single image portrait relighting via explicit multiple reflectance channel modeling. *ACM Transactions on Graphics (TOG)* 39, 1–13.

- [44] Werther, C., Ferguson, M., Park, K., Kling, T., Chen, C., Wang, Y., 2018. Gender effect on face recognition for a large longitudinal database. arXiv .
- [45] Xiang, J., Zhu, G., 2017. Joint face detection and facial expression recognition with mtcnn, IEEE. pp. 424–427.
- [46] Yang, H., Zhu, H., Wang, Y., Huang, M., Shen, Q., Yang, R., Cao, X., 2020. Facescape: a large-scale high quality 3d face dataset and detailed riggable 3d face prediction, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 601–610.
- [47] Zhao, J., Cheng, Y., Xu, Y., Xiong, L., Li, J., Zhao, F., Jayashree, K., Pranata, S., Shen, S., Xing, J., Yan, S., Feng, J., 2018. Towards pose invariant face recognition in the wild, pp. 2207–2216. doi:10.1109/CVPR.2018.00235.
- [48] Zhou, H., Hadap, S., Sunkavalli, K., Jacobs, D.W., 2019. Deep single-image portrait relighting, pp. 7194–7202.
- [49] Zhu, J.Y., Zheng, W.S., Lu, F., Lai, J.H., 2017. Illumination invariant single face image recognition under heterogeneous lighting condition. Pattern Recognition 66, 313–327. doi:10.1016/j.patcog.2016.12.029.
- [50] Zou, X., Kittler, J., Messer, K., 2007. Illumination invariant face recognition: A survey, IEEE. pp. 1–8.