

Face Forgery Detection by 3D Decomposition

Xiangyu Zhu^{1,2*} Hao Wang^{1*} Hongyan Fei³ Zhen Lei^{1,2,4†} Stan Z. Li⁵

¹CBSR & NLPR, Institute of Automation, Chinese Academy of Sciences

²School of Artificial Intelligence, University of Chinese Academy of Sciences

³School of Automation and Electrical Engineering, University of Science and Technology Beijing

⁴Centre for Artificial Intelligence and Robotics, Hong Kong Institute of Science & Innovation, Chinese Academy of Sciences

⁵School of Engineering, Westlake University

{xiangyu.zhu, zlei, szli}@nlpr.ia.ac.cn {haowang7308, hongyanfei0420}@gmail.com

Abstract

Detecting digital face manipulation has attracted extensive attention due to fake media's potential harms to the public. However, recent advances have been able to reduce the forgery signals to a low magnitude. Decomposition, which reversibly decomposes an image into several constituent elements, is a promising way to highlight the hidden forgery details. In this paper, we consider a face image as the production of the intervention of the underlying 3D geometry and the lighting environment, and decompose it in a computer graphics view. Specifically, by disentangling the face image into 3D shape, common texture, identity texture, ambient light, and direct light, we find the devil lies in the direct light and the identity texture. Based on this observation, we propose to utilize facial detail, which is the combination of direct light and identity texture, as the clue to detect the subtle forgery patterns. Besides, we highlight the manipulated region with a supervised attention mechanism and introduce a two-stream structure to exploit both face image and facial detail together as a multi-modality task. Extensive experiments indicate the effectiveness of the extra features extracted from the facial detail, and our method achieves the state-of-the-art performance.

1. Introduction

While earlier seamless face manipulation has amazed the public broadly, there has been a constant concern about the potential abuse of relevant techniques. In particular, the recent *DeepFake* [18] initiated the widespread public discussion among the potential harmful consequence [58] and feasible detection solutions of counterfeit facial media [6].

In this work, we are dedicated to detecting the manip-

*Equal contribution.

†Corresponding author.

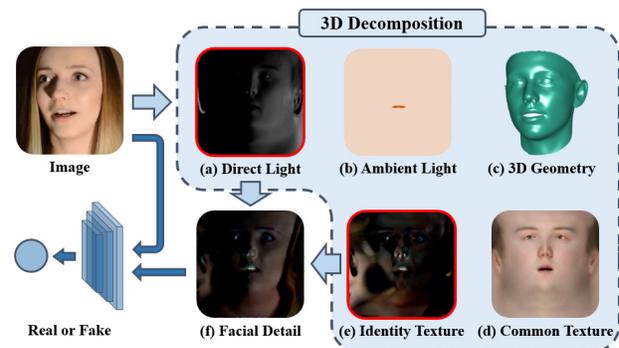


Figure 1. In computer graphics, a face image can be decomposed into direct light, ambient light, 3D geometry, common texture and identity texture. We find critical clues in direct light and identity texture, and merge them as the facial detail for forgery detection.

ulation on facial identity and expression, related to the very popular *DeepFakes* (DF) [18], *Face2Face* (F2F) [57], *FaceSwap* (FS) [34] and *NeuralTextures* (NT) [55], which perform the state-of-the-art face manipulation, making it extremely tough to reveal the sophisticated counterfeit flaws from the image view only [46]. This situation stimulates researchers to shift their attention to extracting forgery evidence from other aspects besides the original RGB image.

Previous work [66, 14, 63, 46] has discovered that the signals in specific frequency ranges are replaced by particular patterns during manipulation and proposes to detect forgery by signal decomposition. The assumption is that, by disentangling the face image, we can find more critical clues for forgery detection from the constituent elements, which are overlooked or hard to be forged by the manipulation methods, whose loss function mainly constrains replication pixel values. For example, Zhang *et al.* [66] identify the unique replications of spectra in the frequency domain due

to the up-sampling process. Chen *et al.* [14] introduce facial semantic segmentation and *Discrete Fourier Transform* (DFT) to extract both spatial- and frequency-domain features, respectively. However, it is difficult to decide which range of signals contains artifacts since images are captured by different devices, under different environments, and even compressed with different algorithms, leading to large frequency distribution bias across datasets. The hand-crafted [52] and learned [46] frequency filters also easily suffer from the generalization problem. Therefore, the crucial problems of this topic lie in how to decompose an image and how to identify reliable constituent elements.

This paper considers the physical decomposition where a face image is the intervention result of its underlying 3D geometry, albedo, and the environment lighting. Specifically, we introduce *3D Morphable Model* (3DMM) [7] and *computer graphics* rendering to simulate the generation of a face image. Under *Lambertian assumption*, we decompose a face image into 5 components (Fig. 1): 3D geometry, common texture, identity texture, ambient light, and direct light. The 3D geometry is the underlying 3D face shape, the common texture is the albedo patterns shared by all the people, the identity texture is the albedo patterns peculiar to this face, the ambient light changes the face color globally, and the direct light generates shading. We introduce the decomposition in Sec. 3.1 in detail.



Figure 2. Samples under strong direct light. The first row is the original faces, the second row is the corresponding fake samples, where evident inconsistency exists in the dim region.

Intuitively, the advanced manipulation methods can well reconstruct 3D geometry, common texture and ambient light since we merely see incompatible facial topology, non-face texture and weird skin color among the massive forged images. Thus, these three elements should be normalized. However, we detect identity texture since it is hard to be simulated due to the rich variations across faces, leading to specific high-frequency artifacts. Besides, we speculate the direct light as another decisive forgery clue with the observation on large artifacts under intense direct light, shown in Fig. 2. By evaluating various compositions among different components, we find that the combination

of **direct light** and **identity texture**, *i.e.*, the **facial detail** in Fig. 1(f), is the best for forgery detection.

When detecting forgery clues with neural networks, we consider the cooperation between face image and facial detail as a multi-modality task and propose a two-stream **Forgery-Detection-with-Facial-Detail** Net (FD²Net). To further highlight the discriminative region, we introduce a supervised Detail-guided Attention mechanism in the network, which employs the facial detail difference between real and fake faces as the objective.

In summary, our contributions are: 1) we introduce 3D decomposition into forgery detection and construct facial detail to amplify subtle artifacts. 2) A two-stream structure FD²Net is proposed to fuse the clues from original images and facial details, where a supervised attention module is introduced to highlight the discriminative region. 3) Compared with the other state-of-the-art detection proposals, our method achieves remarkable elevation on both detection performance and generalization ability.

2. Related Work

Digital face manipulation techniques There has been extensive research on face manipulation. Traditional methods require sophisticated editing tools, domain expertise, and time-consuming processes [59, 61]. Classical CG-based methods explore the correspondence among face geometry and facilitate partial/entire manipulation [56, 57, 53, 33]. Recent deep learning (DL)-based methods, especially with GAN, have demonstrated their power on image synthesis, which promotes both face swapping and synthesis of entire fake images, making it easier to be acquired by the public. While the advanced manipulation techniques based on CG and DL facilitate digital face manipulation remarkably, they exacerbate the difficulty for humans to distinguish manipulated faces from the genuine [50].

General manipulation detection method Abiding concern on potential adverse impacts promotes considerable research on digital manipulation detection. Early methods exploit various artifacts for general image forgery detection [45, 41, 51, 21, 2], and recent advances in *Convolutional Neural Network* (CNN) [35] and DL promote the learning-based methods [8, 3, 31, 68], not only revealing innovative inter-domain artifacts of the generative model [42, 4, 66], but also substantiating the effectiveness of simple supervised classifiers [62]. In particular, facial forgery detection has attracted considerable attention recently, which stimulates massive study according to various forgery techniques [23, 29, 12, 47, 37, 25, 30]. There is a large portion of methods discussing manipulation evidence among low-/high-level features. Raja *et al.* [48] concern the fusion of feature maps of the first fully connected layer from pre-trained CNNs. Zhou *et al.* [67] explore steganalysis features and propose to learn both tampering artifacts and

local noise residual features. Afchar *et al.* [1] exploit the mesoscopic properties of images, arguing that the eyes and mouth influence the detection of *Deepfake* significantly. Liu *et al.* [38] argue the effectiveness and robustness of global/large texture represented by the Gram matrix. There are also methods of transferring images to the frequency domain to explore other forgery evidence [52, 46]. However, previous texture-based methods extract facial features based on pixel-level images, *i.e.*, merely concentrating on exploring manipulation trace among face appearance. Recently, Chen *et al.* [13] highlight the credit of geometry information delivered by 3D shape and achieve remarkable performance in spoof detection, inspiring us to fuse shape and texture clues to find fake faces.

Lighting-based detection There is also research focusing on detecting forgery evidence considering the lighting condition. De Carvalho *et al.* [17] spot forgery evidence from the inconsistency among the 2D illuminant maps of various image segments. Peng *et al.* [44] propose an optimized solution to estimate the 3D lighting environment. However, these methods require comparison among at least two faces in one image, which is problematic in more common scenarios where only one face is in the image.

3. Manipulation Detection with Facial Detail

This paper regards face forgery detection beyond a purely end-to-end binary classification problem. We decompose a face image reversibly into several 3D descriptors, *i.e.*, 3D shape, common texture, identity texture, ambient light and direct light, and explore how these descriptors contribute to the final label, investigating the best combination among them for forgery detection.

3.1. 3D Decomposition

In computer graphics, a face image is generated by:

$$\mathbf{I}_{syn} = Z\text{-Buffer}(\mathbf{S}, \mathbf{C}), \quad (1)$$

where \mathbf{S} is the 3D face mesh, as shown in Fig. 1(c), and \mathbf{C} is the RGB of each vertex in \mathbf{S} . Under the *Lambertian assumption*, the RGB of i th vertex is:

$$\mathbf{C}_i = \mathbf{Amb} * \mathbf{T}_i + \langle \mathbf{n}_i, \mathbf{l} \rangle \cdot \mathbf{Dir} * \mathbf{T}_i, \quad (2)$$

where the facial texture $\mathbf{T}_i = [R_i, G_i, B_i]^T$ is the albedo of the i th vertex, $\mathbf{Amb} = \text{diag}(R_{amb}, G_{amb}, B_{amb})$ is the color of the ambient light, as shown in Fig. 1(b), \mathbf{n}_i is the vertex normal originating from the 3D mesh, \mathbf{l} is the light direction, and $\mathbf{Dir} = \text{diag}(R_{dir}, G_{dir}, B_{dir})$ is the color of the direct light, as shown in Fig. 1(a).

Then, we assume the facial texture \mathbf{T} as the composition of common texture and identity texture, where the common texture \mathbf{T}_{com} is the texture patterns shared by all the people, as shown in Fig. 1(d), and the identity texture \mathbf{T}_{id} is the

discriminative fine-grained texture containing one’s identity information, as shown in Fig. 1(e). In this paper, we model the common texture by the PCA texture model in Basel Face Model (BFM) [43], and calculate the residual between \mathbf{T}_{com} and \mathbf{T} as the identity texture:

$$\mathbf{T} = \bar{\mathbf{T}} + \mathbf{B}\beta + \mathbf{T}_{id}, \quad (3)$$

where $\bar{\mathbf{T}}$ is the mean texture, \mathbf{B} is the principle axes of the PCA texture model, and β is the common texture parameter. Based on these models, any face images can be decomposed by a series of model parameters: $[\mathbf{S}, \mathbf{Amb}, \mathbf{Dir}, \beta, \mathbf{T}_{id}]$, which can be obtained by optimizing the following loss:

$$\arg \min_{\mathbf{S}, \mathbf{Amb}, \mathbf{Dir}, \beta, \mathbf{T}_{id}} \|\mathbf{I} - \mathbf{I}_{syn}(\mathbf{S}, \mathbf{Amb}, \mathbf{Dir}, \beta, \mathbf{T}_{id})\|, \quad (4)$$

where \mathbf{I} is the input face image. After 3D decomposition, the following problems are whether each component contains forgery clues and how to combine them regarding the real/fake label. Firstly, inspired by the previous discussions on high-frequency features beneath pixel-level texture [17, 36], we regard identity texture as a critical forgery clue and remove the topsoil facial texture, *i.e.*, the ambient light and the common texture. Secondly, by observing the fake samples under intensive direct light, as shown in Fig. 2, we can consistently spot artifacts due to the large illumination difference between the source and target faces during manipulation. Therefore, we suppose the existence of forgery clues in the direct light. Moreover, we emphasize the normalization of the 3D shape to make the detector concentrate on fine-grained texture.

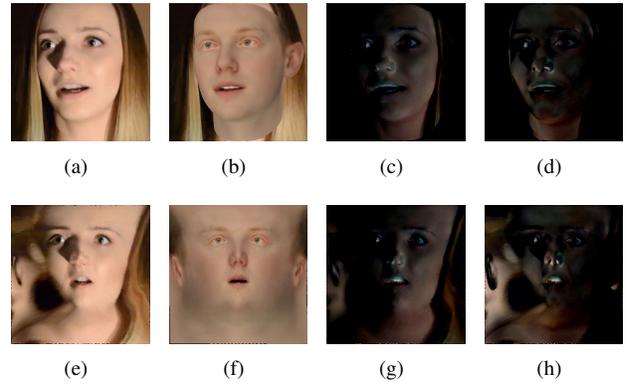


Figure 3. The 8 inputs for forgery detection. (a) Face Image, (b) Ambient Light + Common Texture + 3D Shape, (c) Identity Texture + Direct Light + 3D Shape, (d) Identity Texture + 3D Shape, (e) Face Image w/o Shape, (f) Ambient Light + Common Texture, (g) Identity Texture + Direct Light, (h) Identity Texture.

To verify the suppositions, we conduct a fast ablation study and propose 8 inputs for forgery detection: **img**, **amb+ctex+shape**, **itex+dir+shape**, **itex+shape**, **img w/o shape**, **amb+ctex**, **itex+dir**, **itex**, where amb, dir, itex, ctex

and shape are short for ambient light, direct light, identity texture, common texture and 3D shape, respectively, as shown in Fig. 3, where we warp the image to the UV space to discard the 3D shape. We generate the 8 inputs for all the samples in Faceforensics++ [50] and train a VGG16 for evaluation. The results are shown in Tab. 1.

| input | shape | amb | dir | ctex | itex | AUC |
|-------|-------|-----|-----|------|------|--------------|
| In-a | ✓ | ✓ | ✓ | ✓ | ✓ | 99.13 |
| In-b | ✓ | ✓ | | ✓ | | 50.00 |
| In-c | ✓ | | ✓ | | ✓ | 99.29 |
| In-d | ✓ | | | | ✓ | 99.14 |
| In-e | | ✓ | ✓ | ✓ | ✓ | 98.93 |
| In-f | | ✓ | | ✓ | | 50.00 |
| In-g | | | ✓ | | ✓ | 99.56 |
| In-h | | | | | ✓ | 99.27 |

Table 1. The AUC performance (%) on FFpp [50]. The inputs are the compositions of 5 components, including: 3D shape (shape), ambient light (amb), direct light (dir), common texture (ctex) and identity texture (itex). The examples of In-a to In-h are shown in Fig. 3. The best results are highlighted.

There are several noteworthy results in Tab. 1. Firstly, the poor performances of In-b and In-f indicate that the topsoil facial appearance, *i.e.*, the ambient light and the common texture, are easy to be faked and have little forgery clues. Secondly, by comparing In-c and In-g or comparing In-d and In-h, we find consistent improvements after warping the fine-grained appearance to the UV space. Therefore we suppose normalizing the 3D shape makes CNN concentrate on specific face regions and simplifies the detection task. Thirdly, by comparing In-g and In-h, we find that direct light benefits the forgery detection remarkably. Moreover, we find that many fake samples that In-c identifies but In-d does not are under intense light, verifying our assumption that current manipulation methods cannot simulate direct light properly. In all, the 3D shape, ambient light, and common texture have few forgery patterns but contribute most of the pixel values, and they should be normalized, while the weak signals of direct light and identity texture should be highlighted due to the embedded critical clues. In the following implementation, we use **facial trend** to represent the group of 3D shape, ambient light and common texture, and name the combination of direct light and identity texture as the **facial detail**.

3.2. Facial Detail Generation

Based on the analysis in 3D decomposition, we aim to normalize facial trend (the combination of 3D shape, ambient light and common texture) and highlight facial detail (the combination of direct light and identity texture). A trivial method is optimizing all the parameters together as

Equ. (4) in an analysis-by-synthesis [7] manner, but it costs too much computation. Thus, we propose an approximation to expedite the generation of facial detail for fast inference. We begin with the real-time generation of the 3D shape \mathbf{S} by the state-of-the-art 3DDFA [69, 26, 27], which achieves 0.27ms per image for 3D reconstruction on average. Then, we keep the 3D shape \mathbf{S} and get the ambient and direct light by the spherical harmonic reflectance on the mean texture. With a specialized module, we reduce the average computation time per image from 0.86s to 0.44s compared with the original analysis-by-synthesis method as Equ. (4).

The spherical harmonics [65], $\mathbf{H} = [\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_9]$, are a set of functions that form an orthonormal basis to represent the brightness changes due to illuminations, which can be achieved by the combination of its first nine components [49, 5]. Then, the face appearance under arbitrary illumination can be represented by the linear combination $(\mathbf{H}\boldsymbol{\gamma}) \cdot \mathbf{T}$, where \mathbf{T} is the facial texture (vertex albedo), $\boldsymbol{\gamma} = [\gamma_1, \gamma_2, \dots, \gamma_9]$ is the 9-dimensional reflectance parameters and \cdot is the dot product. We consider $\gamma_1 \cdot \mathbf{h}_1$ as the ambient light and $[\gamma_2 \cdot \mathbf{h}_2, \dots, \gamma_9 \cdot \mathbf{h}_9]$ as the direct light.

In our implementation, we degrade the \mathbf{T} to the mean texture $\bar{\mathbf{T}}$ for fast inference and get $\boldsymbol{\gamma}$ from the least squares solution of the following equation:

$$\mathbf{I}(\mathbf{S}) = (\mathbf{H}\boldsymbol{\gamma}) \cdot \bar{\mathbf{T}}, \quad (5)$$

where $\mathbf{I}(\mathbf{S})$ are the pixels at vertex positions. Based on the harmonic reflectance parameters, we further get the common texture by the following linear equation:

$$\mathbf{I}(\mathbf{S}) = (\mathbf{H}\boldsymbol{\gamma}) \cdot * (\bar{\mathbf{T}} + \mathbf{B}\boldsymbol{\beta}), \quad (6)$$

where $\bar{\mathbf{T}}$ and \mathbf{B} are from the PCA texture model, $\boldsymbol{\gamma}$ is the reflectance parameters estimated in Equ. (5) and $\boldsymbol{\beta}$ is the common texture parameters. Finally, the facial detail is:

$$\mathbf{FD} = UV(\mathbf{I} - (\mathbf{h}_1\gamma_1) \cdot * (\bar{\mathbf{T}} + \mathbf{B}\boldsymbol{\beta}), \mathbf{S}), \quad (7)$$

where \mathbf{FD} is the facial detail and $UV(\mathbf{I}, \mathbf{S})$ is the UV warping that transfers image pixels in \mathbf{I} to the UV space by the constraints of 3D mesh \mathbf{S} . We suppose that the facial detail highlights the forgery patterns for the forged image, making it more suitable for the input of manipulation detection neural networks.

4. FD²Net

We present our **Forgery-Detection-with-Facial-Detail** Net (FD²Net) briefly in Fig. 4, where we adopt Xception-Net [15] as the backbone, considering its state-of-the-art performance on *Faceforensics++* (FFpp) [50], to evaluate the inner-dataset detection performance. Then, we merge the two-stream structure and the attention mechanism into it to explore the potential enhancement assisted by the correlation among the multi-modality data and the location of forgery clues.

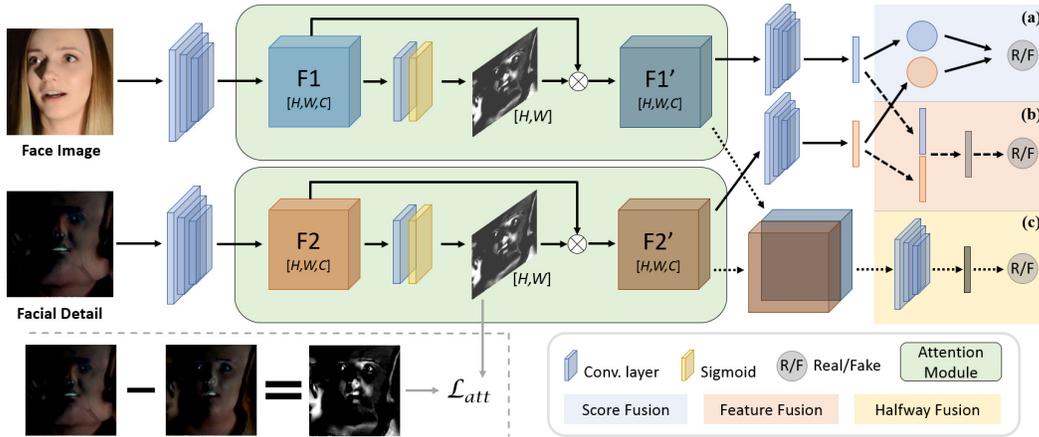


Figure 4. The overview of the **Forgery-Detection-with-Facial-Detail Net (FD²Net)**. We introduce a two-stream structure to combine the clues from original images and facial details. The results of the two streams are fused by three methods: (a) score fusion (SF), (b) feature fusion (FF) and (c) halfway fusion (HF). Besides, we insert a detail-guided attention module, which is supervised by the facial detail difference, in the middle of the backbone network.

4.1. Multi-modality Fusion with Two-stream

Although the facial detail highlights the critical clues in fine-grained texture and shading, it may warp specific forgery patterns and miss external face regions. Therefore, we consider facial detail and original face image as complementary clues and characterize forgery detection as a multi-modality task, regarding facial detail and pixel-level face image as two different modalities. To be specific, we adopt a two-stream architecture to study the combination of these two modalities, where each stream is equipped with the XceptionNet [15] to detect face image and facial detail separately. The classifier performs cooperated decisions at the end regarding two streams’ joint representations.

We evaluate three fusion of two modalities’ representations, shown in Fig. 4. Firstly, we adopt the score fusion (SF), which performs real/fake classification in each stream and adding their confidences as the final score. The result is real if the score is larger than a threshold and fake otherwise, shown in Fig. 4(a). Secondly, we implement the feature fusion (FF), where each stream ends with a fully-connected (FC) layer, and their features are concatenated to a one-dimension vector and transferred by another FC layer to make the final decision, shown in Fig. 4(b). Finally, we design the halfway fusion (HF) to concatenate the intermediate 2D feature maps for further single-stream processing, shown in Fig. 4(c). To be specific, we convolve the two inputs by the first half backbone, *i.e.*, before the 7th block of XceptionNet [15], and stack their 2D outputs as the feature map afterward. Then we adopt the last half to process the stacked feature in a single stream to make the final decision. In our experiments, we find HF performs the best with

smaller parameter size, which may benefit from preserving spatial information when fusing the local features.

4.2. Detail-guided Attention

Extensive tasks have adopted the attention mechanism to enhance forgery detection [16]. The embedded attention module exploits more distinguishing characteristics by positioning the plausible manipulated region, and also strengthens the explainability of the classifiers [16, 9]. Unlike the previous methods, which either need the ground-truth forged regions or adaptively learn the attention map by the real/fake labels [16], we supervise our attention map by the facial detail difference between fake and real faces.

Generally, an attention map \mathbf{M}_{att} is constructed from an intermediate feature map \mathbf{F} by a small regression network $\mathbf{M}_{att} = \mathcal{N}(\mathbf{F}, \theta_{att})$ with θ_{att} as its parameters. Then the intermediate feature is refined by the attention map $\mathbf{F}' = \mathbf{F} \otimes \text{Sigmoid}(\mathbf{M}_{att})$, where \otimes denotes element-wise multiplication. In this work, we propose a novel approach to train the attention network θ_{att} . When constructing the batches during network training, two images are selected for each sample, one real \mathbf{I}_{real} and one fake \mathbf{I}_{fake} . The absolute of the grayscale facial detail difference is taken as a weak supervision of the attention module:

$$\mathcal{L}_{att} = \|\mathcal{N}(\mathbf{F}, \theta_{att}) - |FD(\mathbf{I}_{real}) - FD(\mathbf{I}_{fake})|\|, \quad (8)$$

where $FD(\cdot)$ is facial detail extraction, and the total loss is:

$$\mathcal{L} = \mathcal{L}_{cls} + \mathcal{L}_{att}, \quad (9)$$

where \mathcal{L}_{cls} is the cross entropy loss performing real/fake classification. During test time, we fix the network structure and parameters to infer the attention map.

| Structure | FFpp | | | DFD | | | DFDC | | |
|--------------------|--------------|--------------|-------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | AP | AUC | EER | AP | AUC | EER | AP | AUC | EER |
| Img | 99.44 | 99.31 | 5.39 | 88.07 | 65.57 | 38.38 | 85.60 | 62.17 | 39.99 |
| Detail | 99.40 | 99.12 | 5.51 | 87.24 | 64.29 | 40.87 | 85.02 | 61.80 | 40.37 |
| Img ($\times 2$) | 99.67 | 99.38 | 5.37 | 89.45 | 74.14 | 34.07 | 86.70 | 63.22 | 38.77 |
| Img+Detail (SF) | 97.84 | 92.91 | 11.07 | 83.71 | 72.82 | 36.44 | 81.91 | 62.10 | 47.32 |
| Img+Detail (FF) | 99.72 | 99.45 | 5.31 | 89.56 | 78.55 | 26.80 | 87.16 | 65.36 | 36.17 |
| Img+Detail (HF) | 99.42 | 98.73 | 5.63 | 89.61 | 78.65 | 26.03 | 87.31 | 66.09 | 35.46 |

Table 2. Test results (%) of the two-stream FD²Net and its variants on FFpp, DFD and DFDC. We adopt *Average Precision* (AP), *Area under the Curve* (AUC) of Receiver Operating Characteristic Curve (ROC), and *Equal Error Rate* (EER) as the metrics in our experiments. The “Img” is the stream detecting original images only. The “Detail” is the stream detecting facial details only. The “Img ($\times 2$)” is the one-stream network on original images but having the same parameter size as the two-stream structure. The SF, FF and HF refer to score fusion, feature fusion and halfway fusion, respectively. The best results are highlighted.

5. Experiments

In this section, we introduce the datasets, experiment setups, extensive experiment results of the ablation studies, and comparison with previous works in sequence.

Training Dataset. *Faceforensics++* (FFpp) [50] is a benchmark dataset released recently for facilitating evaluation among facial manipulation detection methods. There are $1k$ original video sequences, in which 720, 140, 140 videos are used for training, validation and testing, respectively. These original videos are manipulated by four state-of-the-art face manipulation methods, *i.e.*, *DeepFakes* (DF) [18], *Face2Face* (F2F) [57], *FaceSwap* (FS) [34], and *NeuralTextures* (NT) [55]. Besides, the raw video sequences are degraded with different compression rate (0, 23, 40) to simulate the real situation [50]. We select the c23 data, considering the extensive post-processing imposed on the original data before they go public, and sample 100 frames for each video in the experiments.

Test Datasets. We adopt the following datasets for performance and generalization evaluation. 1) the testing set of FFpp as described above. 2) *The DeepFake Detection dataset* (DFD) [20] containing hundreds of original data and thousands of manipulated data, released by Google for promoting research on synthetic video detection. 3) *Deepfake detection challenge dataset* (DFDC) [10] containing over 100k video sequences captured with over 3k paid actors and manipulated videos covering Deepfake, GAN-based, and non-learned methods, released recently for the corresponding Kaggle competition¹ by Facebook AI.

Implementation Details. For the facial detail generation, we construct the 3D face shape by 3DDFA [69, 26, 27], perform UV warping by the UV map in [22], and acquire the common texture by fitting the PCA texture model in Basel Face Model (BFM) [43]. For the neural network, we introduce XceptionNet [15] as the backbone and fuse

the feature map after the 4th block of the middle row of XceptionNet when implementing the halfway fusion structure. The Adam optimizer is utilized for training with weight decay equals to 5×10^{-4} , $\beta_1 = 0.9$, $\beta_2 = 0.999$ and batch size set to 32. The initial learning rate is 10^{-4} , then changed to 5×10^{-5} at epoch 15, to 5×10^{-6} at epoch 23, to 10^{-6} at epoch 28, and to 5×10^{-7} for the rest from epoch 32. An early-stop module controls the training process’s end, terminating the training if the loss on the validation set does not fall for 7 consecutive epochs. In our implementation, the total epoch is about 25.

5.1. Ablation Studies

5.1.1 Analysis of the Two-stream Network

We regard face image and facial detail as two complementary modalities and implement a two-stream network to fuse their clues. To evaluate each modality’s performance and the best fusion manner, we quantitatively evaluate FD²Net without attention module in different variants: one-stream with original images only, one-stream with facial details only, and two-stream fused by SF, FF, and HF, respectively. The results are listed in Tab. 2.

Firstly, the one-stream structure considering face images or facial details achieves similar results, but worse than that of the two-stream structure, especially in cross-data evaluation. Secondly, the deteriorated performance of the score fusion indicates the potential sophistication of multi-modality clues fusion and the inflexibility of the hand-crafted decision function. Nevertheless, the performance on all three datasets improves jointly with the feature fusion, validating the complementarity of the two modalities. Finally, the halfway fusion further promotes the results on DFD and DFDC by the local fusion manner, which makes the fused features correspond to similar receptive fields. We also find that the two-stream structure outperforms the one-stream with double parameters, ruling out the benefit from a larger parameter size on the performance improvement.

¹<https://www.kaggle.com/c/deepfake-detection-challenge>

| Attention on Stream | | FFpp | | | DFD | | | DFDC | | |
|---------------------|------------------|--------------|--------------|-------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Img Attention | Detail Attention | AP | AUC | EER | AP | AUC | EER | AP | AUC | EER |
| | | 99.42 | 98.73 | 5.63 | 89.61 | 78.65 | 26.03 | 87.31 | 66.09 | 35.46 |
| ✓ | | 99.45 | 98.73 | 5.62 | 89.65 | 78.71 | 25.94 | 87.58 | 66.51 | 35.33 |
| | ✓ | 99.47 | 98.74 | 5.51 | 89.37 | 78.69 | 26.01 | 87.56 | 66.48 | 35.33 |
| ✓ | ✓ | 99.48 | 98.76 | 5.59 | 89.84 | 79.08 | 25.18 | 87.93 | 67.70 | 34.91 |
| ✓(unsupervised) | ✓(unsupervised) | 99.44 | 98.68 | 5.88 | 88.55 | 78.37 | 27.46 | 87.02 | 65.46 | 37.23 |

Table 3. The ablation study results (%) on the Detail-guided Attention module in FD²Net. The “Img Attention” and “Detail Attention” refer to the attention module on the image stream and detail stream, respectively. We also explore the performance without the supervised signal by the facial detail difference and present the last row results with “unsupervised” in the bracket. The best results are highlighted.

| S1 | S2 | | FFpp | | | DFD | | | DFDC | | |
|-------|----------|------------|--------------|--------------|-------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Image | Tex Norm | Shape Norm | AP | AUC | EER | AP | AUC | EER | AP | AUC | EER |
| ✓ | | | 99.46 | 99.47 | 4.48 | 88.14 | 72.51 | 36.77 | 85.64 | 62.28 | 39.56 |
| ✓ | | ✓ | 99.57 | 99.59 | 4.30 | 84.06 | 76.09 | 29.11 | 86.13 | 64.43 | 38.40 |
| ✓ | ✓ | | 99.61 | 99.68 | 4.28 | 85.10 | 76.84 | 27.08 | 87.73 | 66.01 | 38.32 |
| ✓ | ✓ | ✓ | 99.48 | 98.76 | 5.59 | 89.84 | 79.08 | 25.18 | 87.93 | 67.70 | 34.91 |

Table 4. The ablation study result (%) of facial detail in FD²Net. The S1 and S2 refer to the first and the second stream in the network. The “Tex Norm” refers to texture normalization and the “Shape Norm” refers to shape normalization. The best results are highlighted.

Besides, we tried EfficientNet [54] as the backbone, considering its remarkable performance in DFDC competition, and our method still outperforms the state-of-the-arts with regard to the average precision (93.82% [9] VS. 95.43%) and robustness on DFDC (87.66% [9] VS. 87.85%).

5.1.2 Analysis of the Detail-guided Attention

To highlight the plausible manipulated region, we introduce the detail-guided attention module between the fourth and fifth block of the middle flow of the backbone. Based on the two-stream network with halfway fusion, we evaluate some of the network variants by separately deploying the attention module on each stream and exploring whether the supervised signal improves its effectiveness in further discussion. Results in Tab. 3 demonstrate the improvements of XceptionNet on all datasets with the additional attention module, either implemented on the image or the detail stream. The network with attention modules on both streams performs the best. Besides, we train the attention module indirectly from the real/fake labels, ignoring the supervised signal by the facial detail difference, and observe a performance drop, indicating the effectiveness of the supervised signals on the forgery detection.

5.1.3 Ablation Study of Facial Detail in FD²Net

Although Tab. 1 has performed an extensive ablation study on each 3D component, we further evaluate facial detail when acting as a complementary clue in the two-stream network. In this section, we decompose the effectiveness of

facial detail into shape normalization and texture normalization, where shape normalization refers to warping the facial pixels to the UV space, *e.g.*, Fig. 3(e), and texture normalization refers to decomposing and removing ambient light and common texture in the pixel values, *e.g.*, Fig. 3(c). We adopt the two-stream network with both halfway fusion and supervised attention module.

In Tab. 4, the first row is a one-stream network which directly detects forgery from original face images without using any facial detail information. Adding the second stream can promote the AUC scores compared with the primary one-stream structure, either implementing shape or texture normalization. Furthermore, the introduction of the facial detail, *i.e.*, adopting the combination of shape and texture normalization, helps the detector achieve the best performance. These progressive improvements validate that the proposed facial detail contributes to the forgery detection and complements the original image.

5.2. Comparison with other methods

Some previous works [32, 19, 40, 36] indicate the potential generalization failure when detecting unseen manipulation methods or datasets. In this section, we compare our method with previous state-of-the-art methods to explore our performance in both scenarios.

Cross-data Evaluation. Following Khodabakhsh *et al.* [32], we quantitatively analyze the generalization ability on unseen data and compare it with other methods, including the primary XceptionNet detection method [50] and the ensemble of EfficientNet’s variants [9]. We train the model

| Model | Training dataset | DFD (HQ) | | | DFDC | | |
|-----------------------------|------------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | | AP | AUC | EER | AP | AUC | EER |
| Xception [50] | FFpp | 88.07 | 65.57 | 38.38 | 85.60 | 62.17 | 39.99 |
| EfficientNetB4 Ensemble [9] | FFpp | 89.35 | 72.82 | 34.86 | 85.71 | 63.03 | 38.86 |
| FD ² Net | FFpp | 89.84 | 79.08 | 25.18 | 87.93 | 67.70 | 34.91 |

Table 5. Performance (%) comparison among previous state-of-the-art methods on the unseen datasets. The best results are highlighted.

| Model | Training data | Acc | |
|----------------------|---------------|--------------|--------------|
| | | F2F | FS |
| MesoInception4 [1] | F2F | 84.56 | 56.71 |
| VA-LogReg [39] | | 83.62 | 59.45 |
| LAE [19] | | 90.34 | 62.51 |
| Multi-task [40] | | 91.27 | 55.04 |
| Face X-ray [36] | | 97.73 | 85.69 |
| Xception + HP Filter | | 97.98 | 57.46 |
| FD ² Net | | 98.22 | 86.54 |

Table 6. Detection accuracy comparison (%) with previous methods on F2F and FS in FFpp. We adopt the HQ (c23) version data from FFpp to discuss the robustness on the unseen manipulation technique. The best results are highlighted.

on FFpp, test it on DFD (HQ) and DFDC following the ablation study’s configuration, and list the results in Tab. 5. The improvement in the generalization on unseen data demonstrates that the additional facial detail enables the detection model to effectively extract more discriminative and general features from fake images, even from a different distribution of the training dataset. It is worth noting that the extraction of facial detail is independent of any forgery data, making it the probable reason for better generalization.

Evaluations on Different Manipulation Methods. Following Li *et al.* [36], we evaluate the robustness of our method on the unseen manipulation methods and compare the performance with previous methods. We introduce the data manipulated by different methods, *i.e.*, Face2Face (F2F) and FaceSwap (FS) under the low compression (c23) from FFpp, and train our approach on F2F and test it on both F2F and FS, taking the correct prediction accuracy as the evaluation metric. The results are listed in Tab. 6. The proposed method achieves 98.22% on F2F and 86.54% on FS, with a significant improvement compared to the current state-of-the-art. The improvements mainly benefit from the highlighted clues extracted from the facial detail and the plausible forged regions indicated by the attention map. In particular, some compared methods also consider complementary information from various modalities. Nguyen *et al.* [40] propose sharing knowledge learned simultaneously from images and videos to enhance the performance of the detection on both data. Li *et al.* [36] explore the estimation

of the blending boundary directly from face image to discover the possibility of decomposing an image into the mixture of two images from different sources. Besides, we also include signal decomposition methods in frequency domains, introducing the three base high-pass filters in the *FAD* stream in [46] to primary XceptionNet (Xception + HP Filter). Unlike these methods, the proposed FD²Net strips the ambient lighting and the common appearance with 3D decomposition, exploiting the personalized ambient-free facial detail to extract more robust discriminative features.

6. Conclusion and future work

This paper proposes a novel face forgery detection method by the 3D decomposition of the face image. With the 3D shape, common texture, identity texture, ambient light, and direct light disentangled from the face image, we find critical forgery clues in direct light and identity texture, and propose the facial detail, which is constructed by warping image pixels to the UV space and removing the topsoil facial texture, to highlight the subtle forgery patterns. The clues in the facial detail and the original image are fused by a two-stream network, FD²Net, for the final real/fake classification. Meanwhile, an attention module supervised by the facial detail is proposed to highlight the plausible manipulated region. Extensive experiments demonstrate the effectiveness and robustness of the proposed FD²Net on the FaceForencis++ dataset. We intend to address face analysis with 3D decomposition experimentally to improve the model’s interpretability, which can also be adapted to other facial classification problems, including face recognition [28, 11], liveness detection [64, 13], age estimation [24], etc. For example, in age estimation, we can assume fine-grained 3D shape [60] (*e.g.*, wrinkles, loose skin) as an important factor. In general, our work provides a novel direction to explore the forgery clues by analyzing how an image is generated in physics, following the analysis-by-synthesis idea.

Acknowledgement This work was supported in part by the National Key Research & Development Program (No. 2020AAA0140002), Chinese National Natural Science Foundation Projects #61806196, #61876178, #61976229, #61872367, #61961160704.

References

- [1] Darius Afchar, Vincent Nozick, Junichi Yamagishi, and Isao Echizen. Mesonet: a compact facial video forgery detection network. In *2018 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–7. IEEE, 2018. 3, 8
- [2] Shruti Agarwal and Hany Farid. Photo forensics from jpeg dimples. In *2017 IEEE Workshop on Information Forensics and Security (WIFS)*, pages 1–6. IEEE, 2017. 2
- [3] Jawadul H Bappy, Amit K Roy-Chowdhury, Jason Bunk, Lakshmanan Nataraj, and BS Manjunath. Exploiting spatial structure for localizing manipulated image regions. In *Proceedings of the IEEE international conference on computer vision*, pages 4970–4979, 2017. 2
- [4] Jawadul H Bappy, Cody Simons, Lakshmanan Nataraj, BS Manjunath, and Amit K Roy-Chowdhury. Hybrid lstm and encoder–decoder architecture for detection of image forgeries. *IEEE Transactions on Image Processing*, 28(7):3286–3300, 2019. 2
- [5] Ronen Basri and David W Jacobs. Lambertian reflectance and linear subspaces. *IEEE transactions on pattern analysis and machine intelligence*, 25(2):218–233, 2003. 4
- [6] BBC Bitesize. Deepfakes: What are they and why would i make one? <https://www.bbc.co.uk/bitesize/articles/zfkwcqt>, 2020. 1
- [7] Volker Blanz and Thomas Vetter. Face recognition based on fitting a 3d morphable model. *IEEE Transactions on pattern analysis and machine intelligence*, 25(9):1063–1074, 2003. 2, 4
- [8] Luca Bondi, Silvia Lameri, David Güera, Paolo Bestagini, Edward J Delp, and Stefano Tubaro. Tampering detection and localization through clustering of camera-based cnn features. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1855–1864. IEEE, 2017. 2
- [9] Nicolò Bonettini, Edoardo Daniele Cannas, Sara Mandelli, Luca Bondi, Paolo Bestagini, and Stefano Tubaro. Video face manipulation detection through ensemble of cnns. *arXiv preprint arXiv:2004.07676*, 2020. 5, 7, 8
- [10] Ben Pflaum Jikuo Lu Russ Howes Menglin Wang Cristian Canton Ferrer Brian Dolhansky, Joanna Bitton. The deepfake detection challenge dataset, 2020. 6
- [11] Dong Cao, Xiangyu Zhu, Xingyu Huang, Jianzhu Guo, and Zhen Lei. Domain balancing: Face recognition on long-tailed domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5671–5679, 2020. 8
- [12] Tiago Carvalho, Fabio A Faria, Helio Pedrini, Ricardo da S Torres, and Anderson Rocha. Illuminant-based transformed spaces for image forensics. *IEEE transactions on information forensics and security*, 11(4):720–733, 2015. 2
- [13] Song Chen, Weixin Li, Hongyu Yang, Di Huang, and Yunhong Wang. 3d face mask anti-spoofing via deep fusion of dynamic texture and shape clues. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)(FG)*, pages 137–144. IEEE Computer Society. 3, 8
- [14] Zehao Chen and Hua Yang. Manipulated face detector: Joint spatial and frequency domain attention network. *arXiv preprint arXiv:2005.02958*, 2020. 1, 2
- [15] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017. 4, 5, 6
- [16] Hao Dang, Feng Liu, Joel Stehouwer, Xiaoming Liu, and Anil K Jain. On the detection of digital face manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5781–5790, 2020. 5
- [17] Tiago José De Carvalho, Christian Riess, Elli Angelopoulou, Helio Pedrini, and Anderson de Rezende Rocha. Exposing digital image forgeries by illumination color classification. *IEEE Transactions on Information Forensics and Security*, 8(7):1182–1194, 2013. 3
- [18] deepfakes. Deepfakes github. <https://github.com/deepfakes/faceswap>, Accessed 2020-08-05. 1, 6
- [19] Mengnan Du, Shiva Pentylala, Yuening Li, and Xia Hu. Towards generalizable forgery detection with locality-aware autoencoder. *arXiv preprint arXiv:1909.05999*, 2019. 7, 8
- [20] Nick Dufour and Andrew Gully. Contributing data to deepfake detection research. <https://ai.googleblog.com/2019/09/contributing-data-to-deepfake-detection.html>, 2019. 6
- [21] Hany Farid. *Photo forensics*. MIT press, 2016. 2
- [22] Yao Feng, Fan Wu, Xiaohu Shao, Yanfeng Wang, and Xi Zhou. Joint 3d face reconstruction and dense alignment with position map regression network. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 534–551, 2018. 6
- [23] Javier Galbally, Sébastien Marcel, and Julian Fierrez. Image quality assessment for fake biometric detection: Application to iris, fingerprint, and face recognition. *IEEE transactions on image processing*, 23(2):710–724, 2013. 2
- [24] Xin Geng, Chao Yin, and Zhi-Hua Zhou. Facial age estimation by learning from label distributions. *IEEE transactions on pattern analysis and machine intelligence*, 35(10):2401–2412, 2013. 8
- [25] David Güera and Edward J Delp. Deepfake video detection using recurrent neural networks. In *2018 15th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–6. IEEE, 2018. 2
- [26] Jianzhu Guo, Xiangyu Zhu, and Zhen Lei. 3ddf. <https://github.com/cleardusk/3DDFA>, 2018. 4, 6
- [27] Jianzhu Guo, Xiangyu Zhu, Yang Yang, Fan Yang, Zhen Lei, and Stan Z Li. Towards fast, accurate and stable 3d dense face alignment. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 4, 6
- [28] Jianzhu Guo, Xiangyu Zhu, Chenxu Zhao, Dong Cao, Zhen Lei, and Stan Z Li. Learning meta face recognition in unseen domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6163–6172, 2020. 8
- [29] Abdenour Hadid, Nicholas Evans, Sébastien Marcel, and Julian Fierrez. Biometrics systems under spoofing attack: an

- evaluation methodology and lessons learned. *IEEE Signal Processing Magazine*, 32(5):20–30, 2015. 2
- [30] Javier Hernandez-Ortega, Julian Fierrez, Aythami Morales, and Javier Galbally. Introduction to face presentation attack detection. In *Handbook of Biometric Anti-Spoofing*, pages 187–206. Springer, 2019. 2
- [31] Minyoung Huh, Andrew Liu, Andrew Owens, and Alexei A Efros. Fighting fake news: Image splice detection via learned self-consistency. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 101–117, 2018. 2
- [32] Ali Khodabakhsh, Raghavendra Ramachandra, Kiran Raja, Pankaj Wasnik, and Christoph Busch. Fake face detection methods: Can they be generalized? In *2018 International Conference of the Biometrics Special Interest Group (BIOSIG)*, pages 1–6. IEEE, 2018. 7
- [33] Hyeonwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Niessner, Patrick Pérez, Christian Richardt, Michael Zollhöfer, and Christian Theobalt. Deep video portraits. *ACM Transactions on Graphics (TOG)*, 37(4):1–14, 2018. 2
- [34] Marek Kowalski. Faceswap github. <https://github.com/MarekKowalski/FaceSwap>, Accessed 2020-08-05. 1, 6
- [35] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 2
- [36] Lingzhi Li, Jianmin Bao, Ting Zhang, Hao Yang, Dong Chen, Fang Wen, and Baining Guo. Face x-ray for more general face forgery detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5001–5010, 2020. 3, 7, 8
- [37] Yuezun Li, Ming-Ching Chang, and Siwei Lyu. In ictu oculi: Exposing ai created fake videos by detecting eye blinking. In *2018 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–7. IEEE, 2018. 2
- [38] Zhengzhe Liu, Xiaojuan Qi, and Philip HS Torr. Global texture enhancement for fake face detection in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8060–8069, 2020. 3
- [39] Falko Matern, Christian Riess, and Marc Stamminger. Exploiting visual artifacts to expose deepfakes and face manipulations. In *2019 IEEE Winter Applications of Computer Vision Workshops (WACVW)*, pages 83–92. IEEE, 2019. 8
- [40] Huy H Nguyen, Fuming Fang, Junichi Yamagishi, and Isao Echizen. Multi-task learning for detecting and segmenting manipulated facial images and videos. *arXiv preprint arXiv:1906.06876*, 2019. 7, 8
- [41] James F O’Brien and Hany Farid. Exposing photo manipulation with inconsistent reflections. *ACM Trans. Graph.*, 31(1):4–1, 2012. 2
- [42] Augustus Odena, Vincent Dumoulin, and Chris Olah. Deconvolution and checkerboard artifacts. *Distill*, 1(10):e3, 2016. 2
- [43] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3D face model for pose and illumination invariant face recognition. In *Advanced Video and Signal Based Surveillance, 2009. AVSS’09. Sixth IEEE International Conference on*, pages 296–301. IEEE, 2009. 3, 6
- [44] Bo Peng, Wei Wang, Jing Dong, and Tieniu Tan. Optimized 3d lighting environment estimation for image forgery detection. *IEEE Transactions on Information Forensics and Security*, 12(2):479–494, 2016. 3
- [45] Alin C Popescu and Hany Farid. Exposing digital forgeries by detecting traces of resampling. *IEEE Transactions on signal processing*, 53(2):758–767, 2005. 2
- [46] Yuyang Qian, Guojun Yin, Lu Sheng, Zixuan Chen, and Jing Shao. Thinking in frequency: Face forgery detection by mining frequency-aware clues. In *European Conference on Computer Vision*, pages 86–103. Springer, 2020. 1, 2, 3, 8
- [47] Nicolas Rahmouni, Vincent Nozick, Junichi Yamagishi, and Isao Echizen. Distinguishing computer graphics from natural images using convolution neural networks. In *2017 IEEE Workshop on Information Forensics and Security (WIFS)*, pages 1–6. IEEE, 2017. 2
- [48] Kiran Raja, Sushma Venkatesh, RB Christoph Busch, et al. Transferable deep-cnn features for detecting digital and print-scanned morphed face images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 10–18, 2017. 2
- [49] Ravi Ramamoorthi. Analytic pca construction for theoretical analysis of lighting variability in images of a lambertian object. *IEEE transactions on pattern analysis and machine intelligence*, 24(10):1322–1333, 2002. 4
- [50] Andreas Rossler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. Faceforensics++: Learning to detect manipulated facial images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1–11, 2019. 2, 4, 6, 7, 8
- [51] Husrev Taha Sencar and Nasir Memon. Digital image forensics. *Counter-forensics: attacking image forensics*, pages 327–366, 2013. 2
- [52] José A Stuchi, Marcus A Angeloni, Rodrigo F Pereira, Levy Boccato, Guilherme Folego, Paulo VS Prado, and Romis R-F Attux. Improving image classification with frequency domain layers for feature extraction. In *2017 IEEE 27th International Workshop on Machine Learning for Signal Processing (MLSP)*, pages 1–6. IEEE, 2017. 2, 3
- [53] Supasorn Suwajanakorn, Steven M Seitz, and Ira Kemelmacher-Shlizerman. Synthesizing obama: learning lip sync from audio. *ACM Transactions on Graphics (TOG)*, 36(4):1–13, 2017. 2
- [54] Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019. 7
- [55] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019. 1, 6
- [56] Justus Thies, Michael Zollhöfer, Matthias Nießner, Levi Valgaerts, Marc Stamminger, and Christian Theobalt. Real-time expression transfer for facial reenactment. *ACM Trans. Graph.*, 34(6):183–1, 2015. 2

- [57] Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Face2face: Real-time face capture and reenactment of rgb videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2387–2395, 2016. [1](#), [2](#), [6](#)
- [58] Daniel Thomas. Deepfakes: A threat to democracy or just a bit of fun? <https://www.bbc.com/news/business-51204954>, 2020. [1](#)
- [59] Ruben Tolosana, Ruben Vera-Rodriguez, Julian Fierrez, Aythami Morales, and Javier Ortega-Garcia. Deepfakes and beyond: A survey of face manipulation and fake detection. *arXiv preprint arXiv:2001.00179*, 2020. [2](#)
- [60] Anh Tuan Tran, Tal Hassner, Iacopo Masi, Eran Paz, Yuval Nirkin, and Gérard G Medioni. Extreme 3d face reconstruction: Seeing through occlusions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3935–3944, 2018. [8](#)
- [61] Luisa Verdoliva. Media forensics and deepfakes: an overview. *arXiv preprint arXiv:2001.06564*, 2020. [2](#)
- [62] Sheng-Yu Wang, Oliver Wang, Andrew Owens, Richard Zhang, and Alexei A Efros. Detecting photoshopped faces by scripting photoshop. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 10072–10081, 2019. [2](#)
- [63] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. Cnn-generated images are surprisingly easy to spot... for now. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 7, 2020. [1](#)
- [64] Yan Wang, Song Chen, Weixin Li, Di Huang, and Yuhong Wang. Face anti-spoofing to 3d masks by combining texture and geometry features. In *Chinese Conference on Biometric Recognition*, pages 399–408. Springer, 2018. [8](#)
- [65] Lei Zhang and Dimitris Samaras. Face recognition from a single training image under arbitrary unknown lighting using spherical harmonics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(3):351–363, 2006. [4](#)
- [66] Xu Zhang, Svebor Karaman, and Shih-Fu Chang. Detecting and simulating artifacts in gan fake images. In *2019 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–6. IEEE, 2019. [1](#), [2](#)
- [67] Peng Zhou, Xintong Han, Vlad I Morariu, and Larry S Davis. Two-stream neural networks for tampered face detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1831–1839. IEEE, 2017. [2](#)
- [68] Peng Zhou, Xintong Han, Vlad I Morariu, and Larry S Davis. Learning rich features for image manipulation detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1053–1061, 2018. [2](#)
- [69] Xiangyu Zhu, Xiaoming Liu, Zhen Lei, and Stan Z Li. Face alignment in full pose range: A 3d total solution. *IEEE transactions on pattern analysis and machine intelligence*, 41(1):78–92, 2017. [4](#), [6](#)