EXPLORING SPECULAR REFLECTION INCONSISTENCY FOR GENERALIZABLE FACE FORGERY DETECTION

Anonymous authors

000

001

002003004

010 011

012

013

014

015

016

017

018

019

021

023

024

025

026

027

028

029

031

033

035

037

038

040

041

042

043

044

046

047

048

051

052

Paper under double-blind review

ABSTRACT

Detecting deepfakes has become increasingly challenging as forgery faces synthesized by AI-generated methods, particularly diffusion models, achieve unprecedented quality and resolution. Existing forgery detection approaches relying on spatial and frequency features demonstrate limited efficacy against high-quality, entirely synthesized forgeries. In this paper, we propose a novel detection method grounded in the observation that facial attributes governed by complex physical laws and multiple parameters are inherently difficult to replicate. Specifically, we focus on illumination, particularly the specular reflection component in the Phong illumination model, which poses the greatest replication challenge due to its parametric complexity and nonlinear formulation. We introduce a fast and accurate face texture estimation method based on Retinex theory to enable precise specular reflection separation. Furthermore, drawing from the mathematical formulation of specular reflection, we posit that forgery evidence manifests not only in the specular reflection itself but also in its relationship with corresponding face texture and direct light. To address this issue, we design the Specular-Reflection-Inconsistency-Network (SRI-Net), incorporating a two-stage cross-attention mechanism to capture these correlations and integrate specular reflection related features with image features for robust forgery detection. Experimental results demonstrate that our method achieves superior performance on both traditional deepfake datasets and generative deepfake datasets, particularly those containing diffusion-generated forgery faces.

1 Introduction

With the rapid advancement of face manipulation techniques, generating highly realistic forgery face images and videos has become increasingly effortless (Rombach et al., 2022; Chen et al., 2025), and the potential abuse of this technology raises significant security concerns. Particularly with the rise of AI-generation tools such as diffusion models, which enable more sophisticated, realistic, and higher-resolution entire synthesized forgeries, making the challenge of face forgery detection becomes even harder (Kawar et al., 2023; Rosberg et al., 2023; Kim et al., 2025).

Current face forgery detection methods primarily identify forgery evidence by analyzing discriminative features extracted from the spatial (Li et al., 2020a; Cao et al., 2022) or frequency (Tan et al., 2024) domains of face images, as well as leveraging pre-trained features (Cui et al., 2025) from models such as CLIP (Radford et al., 2021). However, these approaches face significant limitations as texture inconsistencies in contemporary forgery faces, particularly those entirely synthesized by generative models, become increasingly subtle. Moreover, forgery evidence from different generation methods may manifest in distinct frequency bands, while pre-trained features lack domain-specific knowledge and interpretability. These factors collectively contribute to the declining detection performance of existing methods when confronted with sophisticated synthetic faces.

In this paper, we propose an approach to localize forgery evidence from the perspective of face generation. Specifically, we posit that despite the remarkable realism achieved by state-of-the-art forgery techniques, they remain constrained by a fundamental principle: attributes governed by more estimated parameters and more complex physical laws are inherently more challenging to replicate accurately. Illumination represents a particularly complex physical phenomenon, as it is simultaneously governed by 2D local texture statistics, 3D global environmental conditions, and

056

058

060 061 062

063 064

065

066

067

068 069

071

072

073

074 075

076

077

079

081

083

084

085

087

880

090

092

093

094 095

096

098

099

100

101 102 103

104 105

107

Figure 1: The visualization of (a) Spatial-based face forgery detection methods, (b) Frequency-based face forgery detection methods and (c) Our proposed Specular-Reflection-Inconsistency-Network (SRI-Net). SRI-Net analyzes that specular reflection is more difficult to replicate based on the mathematical form of the general Phong illumination model and contains generalizable forgery evidence.

illumination model constraints, making it a promising avenue for investigation. As illustrated in Fig. 1 (c), under the Phong illumination model, face illumination comprises ambient light, direct light, and specular reflection. The mathematical representation of the specular reflection component involves more parameters and exhibits stronger nonlinearity, making it inherently more difficult to accurately replicate.

Based on this analysis, we hypothesize that specular reflection contains more generalizable forgery evidence. In this paper, we introduce the 3D Morphable Model (3DMM) (Blanz & Vetter, 2003) and computer graphics rendering algorithms to achieve illumination/texture separation and specular reflection extraction. However, since the widely used Basel Face Model (BFM) based texture estimation (Paysan et al., 2009) cannot capture fine-grained facial textures, the extracted specular reflection component inevitably contains erroneous details. To address this limitation, we propose a faster and more accurate face texture estimation method based on Retinex theory (Land & Mc-Cann, 1971) to replace the BFM model. Under the constraints of Retinex-based texture, we achieve more precise extraction of specular reflection. Furthermore, according to the analytical expression of specular reflection in the Phong illumination model, its estimation depends on both direct light intensity and reflective material properties. Since human faces share similar skin properties, the relative material characteristics can be approximated through face texture. Consequently, forgery evidence should manifest not only in the specular reflection component itself but also in its relationships with direct light and face texture. To exploit these relationships, we design the Specular-Reflection-Inconsistency-Network (SRI-Net) with a two-stage cross-attention structure to effectively capture the correlations among these attributes. The network subsequently integrates specular reflection related features with image features for final real/fake decision.

Our contributions can be summarized as follows:

- 1) We identify specular reflection as a robust forgery indicator due to its complex multi-parameter formulation, stronger nonlinearity and inherent difficulty to replicate accurately.
- 2) We introduce a Retinex-based face texture estimation method that enables faster and more precise specular reflection extraction.
- 3) We design SRI-Net with two-stage cross-attention to exploit the relationships among specular reflection, face texture, and direct light.
- 4) Extensive experiments demonstrate superior performance on both traditional and generative deepfake datasets.

2 RELATED WORK

2.1 General Face Forgery Detection

Many previous studies have focused on general face forgery detection. From the perspective of feature extraction, these approaches can be broadly categorized into spatial-domain and frequency-

domain approaches. Spatial-domain methods primarily analyze pixel-level forgery evidence and inconsistencies in forgery images. MesoNet (Afchar et al., 2018) and Face X-ray (Li et al., 2020a) focused on mesoscopic properties or boundary discrepancies to distinguish manipulated faces. Some works (Wang et al., 2020; Wang & Deng, 2021; Cao et al., 2022) employed techniques such as adversarial training and reconstruction for enriching training data. Frequency-domain approaches (Masi et al., 2020; Qian et al., 2020; Jeong et al., 2022; Tan et al., 2024) decomposed images into frequency components using Fourier or wavelet transforms to detect forgery traces that are less perceptible in the spatial domain. From the perspective of learning generalized representations, SBI (Shiohara & Yamasaki, 2022) and SLADD (Chen et al., 2022) explicitly minimized domain gaps to improve cross-dataset detection. Some works (Sun et al., 2020; Yan et al., 2024a) employed data augmentation and synthesis in latent space to boost model robustness. Some methods (Yan et al., 2023; Dong et al., 2023) incorporated generalized forgery cues through specialized learning frameworks for more discriminative detection. Some methods (Cui et al., 2025; Sun et al., 2025) leverage pretrained features from models such as CLIP. However, as forgery faces become increasingly realistic and forgery evidence is distributed unevenly across frequency bands, and with new forgery techniques continuously emerging, the effectiveness of existing detection methods declines.

2.2 Illumination-based Detection

Several existing methods (Peng et al., 2016; Zhu et al., 2021) detected forgery evidence by extracting illumination cues from 3D disentangled faces under the Lambertian assumption, demonstrating notable performance. Some works (Hu et al., 2021; Ebihara et al., 2020; Seibold et al., 2018) captured the local specular reflection on the cornea or nose for face forgery detection and face anti-spoofing. However, recent studies (Du et al., 2023; Zhan et al., 2024) revealed that diffusion models can effectively interpret geometric patterns and generate realistic shading estimations, thereby diminishing the reliability of such illumination-based detection approaches.

Nevertheless, the fundamental principle that attributes requiring more estimated parameters remain inherently more challenging to replicate remains unchanged. Given that illumination constitutes a complex physical process, it warrants further exploration. Therefore, we propose extending the Lambertian assumption to the more physically realistic Phong illumination model and precisely extract the complete and precise specular reflection component under 3D physical constraints, capturing the potential forgery evidence within it.

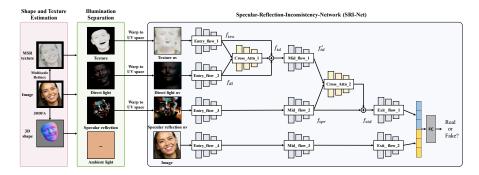


Figure 2: The framework of our proposed face forgery detection method. First, we use 3DDFA to extract 3D shape and propose a fast and accurate Retinex-based method for texture extraction. Next, we employ spherical harmonic model to fit ambient and direct light, extracting specular reflection through a residual based approach under Retinex-based texture constraints. We then propose the Specular-Reflection-Inconsistency-Network (SRI-Net) with a two-stage cross-attention structure to capture correlations among specular reflection, texture, and direct light. Finally, SRI-Net combines these specular reflection related features with image features for final real/fake decision.

3 METHOD

In this section, we first analyze face illumination under the Phong illumination model and demonstrate that specular reflection is more difficult to replicate and contains more generalizable forgery

evidence. Subsequently, we propose a Retinex theory based face texture extraction method to achieve faster and more accurate specular reflection extraction. Finally, we present SRI-Net with a two-stage cross-attention architecture to capture the correlations of specular reflection, texture, and direct light, enabling more comprehensive forgery evidence extraction. The framework of our proposed method is illustrated in Fig. 2.

3.1 FACE ILLUMINATION ANALYSIS

From the perspective of face generation, a face image is composed of 3D structure and color information:

$$\mathbf{I}_{sum} = \mathcal{R}(\mathbf{S}, \mathbf{C}),\tag{1}$$

where \mathcal{R} is the renderer, \mathbf{S} is the 3D shape, and \mathbf{C} represents the color information that combines texture (i.e., skin albedo) and illumination effects for each vertex on the 3D shape. The 3D shape is the easiest to replicate, as there is almost no occurrence of incompatible facial topology in existing forgery faces. While the color information \mathbf{C} integrates both texture and illumination, the illumination component, as a complex physical process, is simultaneously constrained by both 2D local texture statistics, 3D global environmental conditions, and illumination model constraints. Therefore, it warrants further exploration to extract potential forgery evidence from the illumination-related aspects of the color information.

The Phong illumination model is a widely used local illumination model that approximates the interaction of light with surfaces by combining ambient light, direct light and specular reflection components. Ambient light is uniform and illuminates all surfaces equally, which ensures that objects are visible even in areas not directly illuminated. Direct light is scattered across a surface based on its angle to the surface normal. Specular reflection creates shiny highlights, dependent on the surface's reflection of the light. Under Phong assumption, the RGB value for each vertex C_i on S is computed as follows:

$$C_i = \mathbf{Amb} * \mathbf{T}_i + \langle \mathbf{n}_i, \mathbf{l} \rangle \cdot \mathbf{Dir} * \mathbf{T}_i + \langle \mathbf{r}_i, \mathbf{v} \rangle^n \cdot \mathbf{Dir} * \mathbf{T}_i,$$
(2)

where Amb is ambient light, Dir is direct light, T is the texture (i.e., skin albedo), $\langle \mathbf{r}_i, \mathbf{v} \rangle^n \cdot \mathbf{Dir}$ is specular reflection, \langle , \rangle denotes the inner product operation, and * denotes the element-wise multiplication. \mathbf{n}_i is the vertex normal of the 3D shape, 1 is the direct light direction, \mathbf{r}_i is the reflection direction which can be computed by $\mathbf{r}_i = 2\langle \mathbf{n}_i, 1 \rangle \mathbf{n}_i - \mathbf{l}$, \mathbf{v} is the viewer direction, and n is the specular exponent of surface material. It can be observed that among the three types of illuminations, the mathematical formulation of specular reflection is more complex, involving six estimated parameters and a stronger nonlinear representation in exponential form. This makes it inherently more difficult to replicate.

Due to the fact that the Phong illumination model is a general model with clear physical significance, this characteristic where specular reflection is more difficult to replicate can be applied across various forgery methods, including current generative forgery techniques. Therefore, we posit that the specular reflection contains more generalizable forgery evidence.

3.2 SPECULAR REFLECTION EXTRACTION

To extract the specular reflection component, two main steps are involved according to Equ. 2: 1) illumination and texture separation, and 2) illumination components separation.

Illumination and texture separation

Existing methods typically perform illumination and texture separation using an analysis-by-synthesis approach within the 3D Morphable Model (3DMM) (Blanz & Vetter, 2003). The illumination and texture are parameterized through the spherical harmonic model (Zhang & Samaras, 2006) and the PCA texture model of Basel Face Model (BFM) (Paysan et al., 2009), with iterative optimization of the coefficients to obtain the illumination and texture components. In spherical harmonic model, face color under arbitrary illumination can be represented by the linear combination:

$$\mathbf{I}(\mathbf{S}) = (\mathbf{H}\gamma) \cdot \mathbf{T},\tag{3}$$

where $\mathbf{H} = [\mathbf{h}_1, \mathbf{h}_2, ..., \mathbf{h}_n]$ denotes a set of harmonic reflectance functions forming orthonormal bases to model illumination-induced brightness variations, while $\gamma = [\gamma_1, \gamma_2, ..., \gamma_n]$ represents the n-dimensional reflectance parameters requiring estimation. Notably, \mathbf{h}_1 is isotropic and $\mathbf{h}_1 \cdot \gamma_1$

 can represent ambient light, whereas the directional components $[\mathbf{h}_2 \cdot \gamma_2, ..., \mathbf{h}_9 \cdot \gamma_9]$ can represent direct light. Basis with n>9 can be used to fit illumination with more detailed illumination, but the computational cost also increases significantly.

In PCA texture model of BFM, the texture is parameterized in the following form:

$$\mathbf{T} = \overline{\mathbf{T}}_{bfm} + \mathbf{B}\beta,\tag{4}$$

where $\overline{\mathbf{T}}_{bfm}$ is the mean face texture of BFM, \mathbf{B} is the 199 dimensional principle axes, β is the corresponding texture parameter. Based on the above equations, the parameters can be optimized through iterative methods by:

$$Argmin_{\gamma,\beta} \|\mathbf{I} - \mathbf{I}_{syn}(S,\gamma,\beta)\|. \tag{5}$$

where I is the original face image. The estimated $\mathbf{H}\gamma$ and $\overline{\mathbf{T}}_{bfm} + \mathbf{B}\beta$ can represent the separated illumination and texture.

However, the PCA bases **B** are obtained by eigendecomposition of the covariance matrix of 200 3D face scans in BFM, capturing dominant shape variations but failing to reconstruct identity fine-grained texture details due to dimensionality reduction and linear modeling constraints. Therefore, in real-world scenario, the face texture should be formulated as follows:

$$\mathbf{T} = \overline{\mathbf{T}}_{bfm} + \mathbf{B}\beta + \mathbf{T}_{id},\tag{6}$$

where \mathbf{T}_{id} is the identity texture details of corresponding face image. This indicates that the result derived from Equ. 5 contains a residual term \mathbf{T}_{id} when compared to the real texture \mathbf{T} . Furthermore, inaccurate estimation of the texture will further lead to inaccurate estimation of the illumination. Particularly, due to the more complex and fine-grained nature of the specular reflection, it is more significantly affected when constrained by an inaccurate texture. Therefore, it is necessary to propose a new texture estimation method with a smaller residual to the real texture. To address this issue, we propose a Retinex theory (Land & McCann, 1971) based texture extraction method for faster and more accurate estimation of specular reflection. Retinex theory is grounded in the physical model of image formation, which assumes that an observed image I(x,y) is the product of illumination L(x,y) and albedo R(x,y), expressed as $I(x,y) = L(x,y) \cdot R(x,y)$. To separate these components, Retinex employs a logarithmic transformation, converting the multiplicative relationship into an additive one:

$$\log I(x,y) = \log L(x,y) + \log R(x,y),\tag{7}$$

where the illumination L(x,y) is typically estimated using low-pass filtering (e.g., Gaussian filtering) of the input image. This choice is theoretically justified by the assumption that illumination varies spatially smoothly, while albedo contains high-frequency details such as identity texture. By applying a Gaussian filter G_{σ} , the illumination is approximated as $G_{\sigma}(I(x,y))$, and the albedo can be obtained by:

$$\log R(x,y) = \log I(x,y) - \log[G_{\sigma}(I(x,y))], \tag{8}$$

where logarithmic form $\log R(x,y)$ is preferred over linear form R(x,y) to represent the texture in Retinex theory. This is because it avoids nonlinear exponential transformations that could reintroduce illumination artifacts or cause numerical instability. The log-domain processing maintains illumination invariance while better preserving fine surface details, making it more suitable for computational analysis of intrinsic surface properties under varying illumination conditions. Thus, the texture \mathbf{T}_{re} of face image obtained by Retinex can be represented by the following formula:

$$\mathbf{T}_{re} = \log R. \tag{9}$$

Furthermore, to simultaneously address both global and local illumination variations, we employ Multi-Scale Retinex (MSR) to enhance robustness under complex illumination conditions. Specifically, large-scale Gaussian kernels are utilized to capture broad illumination gradients, while small-scale kernels supplement local illumination details, thereby enabling stable texture extraction across diverse illumination scenarios. The formula is presented as follows:

$$\mathbf{T}_{msr} = \frac{1}{N} \sum_{i=1}^{N} \left(\log I - \log \left[G_{\sigma_i}(I) \right] \right), \tag{10}$$

where G_{σ_i} denotes Gaussian kernels with varying scales σ_i , and N represents the number of scales. This multi-scale integration ensures comprehensive illumination normalization while preserving fine-grained texture information. Thus, the illumination and texture separation process can be transformed from Equ. 5 into the following formula:

$$Argmin_{\gamma} \|\mathbf{I} - \mathbf{I}_{syn}(S, \gamma, \mathbf{T}_{msr})\|. \tag{11}$$

Notably, substituting $\mathbf{T} = \overline{\mathbf{T}} + \mathbf{B}\beta$ with $\mathbf{T} = \mathbf{T}_{msr}$ not only makes the texture estimation more accurate but also eliminates the necessity to estimate the β parameter. This approach circumvents the additional computational burden associated with the iterative optimization of both β and γ parameters required for the simultaneous optimization of illumination and texture. As a result, the processing time for single-image illumination and texture separation is reduced from 0.78s to 0.29s, significantly improving efficiency and enabling rapid large-scale data processing.

Illumination components separation

Based on Equ. 11, the spherical harmonic coefficients can be estimated. An intuitive approach is to utilize $\mathbf{h}_1 \cdot \gamma_1$ as the ambient light, $[\mathbf{h}_2 \cdot \gamma_2, ..., \mathbf{h}_9 \cdot \gamma_9]$ as direct light, and higher order basis along with their coefficients to represent the specular reflection. However, this approach faces similar issues to the PCA model of BFM, namely that the specular reflection fitted using spherical harmonic bases still exhibits residuals when compared to the real specular reflection. Additionally, the estimation of higher order spherical harmonic coefficients introduces significant computation cost.

To address this issue, we propose separating the specular reflection through a residual based approach. We utilize the first nine basis functions of the spherical harmonic model to capture coarse illumination variations, which effectively fit the ambient light and direct light while being insensitive to fine-grained specular reflection. Therefore, the separation of specular reflection can be achieved in the following form:

$$SPR = (I - (H\gamma)_{(1-9)} \cdot T_{msr})/T_{msr}, \tag{12}$$

where **SPR** is specular reflection. $(\mathbf{H}\gamma)_{(1-9)}$ represents the first nine spherical harmonic bases along with the coefficients estimated according to Equ. 11.

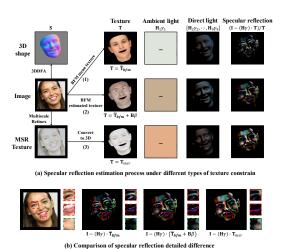


Figure 3: The visualization of (a) Specular reflection estimation process under different types of texture constraints and (b) Comparison of specular reflection detailed difference.

Fig. 3 (a) presents visualization of using (1) the mean texture $\overline{\mathbf{T}}_{bfm}$ and (2) the fitted texture $\overline{\mathbf{T}}_{bfm} + \mathbf{B}\beta$ from BFM and (3) the Retinex-based texture \mathbf{T}_{msr} as texture \mathbf{T} , along with their corresponding obtained specular reflection, where the 3D shape is obtained by 3DDFA (Guo et al., 2020) in real-time. The results show that our proposed \mathbf{T}_{msr} achieves more accurate texture estimation compared to $\overline{\mathbf{T}}_{bfm}$ and $\overline{\mathbf{T}}_{bfm} + \mathbf{B}\beta$, thereby enabling fine-grained removal of identity texture from specular reflection. Fig. 3 (b) presents a comparison of detailed differences in specular reflection. The red bounding boxes demonstrate that specular reflection with \mathbf{T}_{msr} as constraint achieves more accurate

removal of identity texture (e.g., in ocular and labial regions), while the green bounding boxes confirm that the extracted specular reflection better aligns with physical realism.

3.3 SRI-NET

We propose a Specular-Reflection-Inconsistency-Network (SRI-Net) with Xception (Chollet, 2017) as backbone, with its architecture illustrated in Fig. 2. As shown in Equ. 2, the specular reflection intensity is determined by the light direction, 3D shape vertex normal, viewer direction, specular exponent of surface material, direct light intensity, and face texture. Therefore, forgery evidence should manifest not only in the specular reflection itself but also in its correlations with the aforementioned attributes. Given that human faces are all made of skin, the relative material properties can be approximated through face texture.

To capture these correlations, we first flatten the specular reflection, texture, and direct light components into UV space to normalize their directions. Subsequently, we employ a two-stage cross-attention mechanism with residual connections to extract their correlations. The first stage captures the relationship between texture and direct light. Let f_{tex} and f_{dl} denote the flattened feature maps of texture and direct light after the entry flow. The aggregated feature is computed as:

$$f_{td} = \text{Softmax}\left(\frac{f_{tex}f_{dl}^T}{\sqrt{d}}\right)f_{dl} + f_{tex} + f_{dl},\tag{13}$$

where d is the feature dimension, and f_{td} encodes both features and their correlations. The second stage captures the correlations between specular reflection and the texture-direct light features. Let f'_{td} be the processed f_{td} through the middle flow, and f_{spr} be the specular reflection features after passing through both entry and middle flows:

$$f_{std} = \text{Softmax}\left(\frac{f'_{td}f^T_{spr}}{\sqrt{d}}\right) f_{spr} + f_{spr},\tag{14}$$

where f_{std} aggregates the specular reflection features along with the correlations among specular reflection, texture, and direct light.

Meanwhile, since the original image retains the initial forgery evidence, we introduce an image branch to extract subtle information that may be distorted during specular reflection extraction, serving as a complementary component. Finally, we fuse the specular reflection related features with the image features to form the feature vector for final real/fake decision.

4 EXPERIMENTS

In this section, we provide a comprehensive evaluation of SRI-Net, covering various aspects such as datasets, detection performance comparisons, and ablation studies to demonstrate the effectiveness of SRI-Net.

4.1 Datasets

Traditional datasets. We employ the FaceForensics++ (FF++) (Rossler et al., 2019), CelebDF_v1, CelebDF_v2 (Li et al., 2020b), and DeepfakeDetection (DFD) (Google AI, 2019) datasets for evaluation. FF++ is a widely used benchmark dataset designed for evaluating face manipulation detection methods. It contains 1,000 original video sequences, each altered using four state-of-the-art face manipulation techniques: DeepFakes, Face2Face, FaceSwa, and NeuralTextures. To simulate real-world conditions, we utilize the FF++ c23 (light compression) version as the training set in subsequent experiments.

Generative datasets. We employ Diffusion Facial Forgery (DiFF) (Cheng et al., 2024a) and DF40 (Yan et al., 2024b) datasets for evaluation. DiFF dataset contains over 500,000 images synthesized using various generation methods (e.g., Stable Diffusion XL (SDXL), Low-Rank Adaptation (LoRA)) across four categories: Text-to-Image (T2I), Image-to-Image (I2I), Face Swapping (FS) and Face Editing (FE). Additionally, DF40 dataset contains forged data generated within the FF++ and CelebDF domains to ensure method diversity while maintaining consistent data distribution.

Evaluation Metrics. We utilize the Frame-Level Area Under Curve (AUC) metric on generative deepfake datasets, and utilize both frame-level and video-level AUC on traditional deepfake datasets.

4.2 COMPARISON ON TRADITIONAL DATASETS

Table 1: Frame-level AUC (%) performance comparison on traditional Datasets.

Method	Venue	Traditional Datasets				
Wethod	venue	CDF-v1	CDF-v2	DFD	Avg	
FWA (Li & Lyu, 2018)	CVPRW'18	79.0	66.8	74.0	73.3	
CapsuleNet (Nguyen et al., 2019)	ICASSP'19	79.1	74.7	68.4	74.1	
CNN-Aug (Haliassos et al., 2022)	CVPR'20	74.2	70.3	64.6	69.7	
Face X-ray (Li et al., 2020a)	CVPR'20	70.9	67.9	76.6	71.8	
FFD (Dang et al., 2020)	CVPR'20	78.4	74.4	80.2	77.7	
F3Net (Qian et al., 2020)	ECCV'20	77.7	79.8	70.2	75.9	
SPSL (Liu et al., 2021)	CVPR'20	81.5	76.5	81.2	79.7	
SRM (Luo et al., 2021)	CVPR'21	79.3	75.5	81.2	78.7	
CORE (Ni et al., 2022)	CVPRW'22	78.0	74.3	80.2	77.5	
RECCE (Cao et al., 2022)	CVPR'22	76.8	73.2	81.2	77.1	
SBI (Shao et al., 2022)	CVPR'22	-	81.3	77.4	-	
UCF (Yan et al., 2023)	ICCV'23	77.9	75.3	80.7	78.0	
ED (Ba et al., 2024)	AAAI'24	81.8	86.4	-	-	
ProDet (Cheng et al., 2024b)	NIPS'24	90.9	84.2	84.8	86.6	
LSDA (Yan et al., 2024a)	CVPR'24	86.7	83.0	88.0	85.9	
Fada (Cui et al., 2025)	CVPR'25	-	83.7	-	-	
FIA-USA (Ma et al., 2025)	ARXIV'25	90.1	86.7	82.1	86.3	
SRI-Net (Ours)	-	91.3	87.5	89.3	89.4	

Tab. 1 and Tab. 2 present the frame-level and video-level AUC comparisons on traditional datasets, where video-level scores are computed by averaging the frame-level scores across all frames. SRI-Net achieves state-of-the-art performance under both evaluation protocols. Specifically, it attains the highest AUC scores across all three datasets at the frame level, demonstrating its effectiveness in capturing subtle forgery evidence through explicit modeling of correlations among specular reflection, texture, and direct light. At the video level, SRI-Net maintains its superior performance despite relying solely on frame-wise score aggregation without exploiting temporal dependencies.

Table 2: Video-level AUC(%) performance comparison on traditional Datasets.

Method	Traditional Datasets			
Method	CDF-v2	DFD		
Xception	81.6	89.6		
EffiNet-B4	80.8	86.2		
RECCE	82.3	89.1		
F3-Net	78.9	84.4		
SBI	90.6	88.2		
UCF	83.7	86.7		
FIA-USA	94.1	-		
SRI-Net (Ours)	95.5	93.1		

Table 3: Frame-level AUC (%) performance comparison on DF40 Dataset.

Method	Venue	DF40 Dataset						
Method	venue -	uniface	e4s	facedancer	fsgan	inswap	simswap	Avg
RECCE	CVPR'22	84.2	65.2	78.3	88.4	79.5	73.0	78.1
SBI	CVPR'22	64.4	69.0	44.7	87.9	63.3	56.8	64.4
CORE	CVPRW'22	81.7	63.4	71.7	91.1	79.4	69.3	76.1
IID	CVPR'23	79.5	71.0	79.0	86.4	74.4	64.0	75.7
UCF	ICCV'23	78.7	69.2	80.0	88.1	76.8	64.9	76.3
LSDA	CVPR'24	85.4	68.4	75.9	83.2	81.0	72.7	77.8
CDFA	ECCV'24	76.5	67.4	75.4	84.8	72.0	76.1	75.4
ProgressiveDet	NIPS'24	84.5	71.0	73.6	86.5	78.8	77.8	78.7
FIA-USA	ARXIV'25	91.8	87.5	83.0	86.3	87.4	91.0	87.8
SRI-Net (Ours)	-	92.0	89.4	95.3	94.2	91.1	83.3	90.9

4.3 Comparison on Generative Datasets

Tab. 3 and Tab. 4 present frame-level AUC comparisons on the DF40 and DiFF datasets. For the six representative faceswap subsets in DF40 dataset, SRI-Net achieves a substantial improvement in av-

erage AUC from 87.8% to 90.9%. Similarly, on DiFF dataset, our method consistently outperforms existing approaches across all four subsets.

Notably, while conventional methods often experience performance degradation on generative datasets due to the reduced texture inconsistencies, SRI-Net maintains robust detection performance. This resilience can be attributed to the fact that specular reflection is inherently difficult to replicate accurately, and modeling their correlations with multiple facial attributes based on well established physical principles. This physics grounded approach enables SRI-Net to capture generalizable forgery evidence that remains present even in high quality synthetic forgery faces.

Table 4: Frame-level AUC (%) performance comparison on DiFF Dataset.

Method	DiFF Dataset					
Michied -	T2I	I2I	FS	FE	Avg	
Xception	62.4	56.8	86.0	58.6	66.0	
EffiNet-B4	74.1	57.3	82.1	57.2	67.6	
F3-Net	66.9	67.7	81.0	60.6	69.1	
SBI	80.2	80.4	85.1	68.8	78.6	
FIA-USA	86.1	85.0	89.4	72.7	83.3	
SRI-Net (Ours)	88.7	85.5	92.2	80.8	86.8	

4.4 ABLATION STUDIES

In this section, we conduct ablation studies to analyze the SRI-Net, the experimental results are presented in Tab. 5.

Effectiveness of each branch. Rows 1-4 in Tab. 5 evaluate individual branch contributions. The combination of image (Img) and specular reflection (SPR) branches achieves superior performance, as the original image preserves undistorted forgery evidence while the SPR branch isolates difficult to replicate specular reflection that contains generalizable forgery evidence. Adding the texture/direct light (Tex/Dir) branch further improves results by leveraging the Phong illumination model's physical principles.

Table 5: Ablation study results on AUC (%) of SRI-Net.

Method Settings	Datasets				
Method Settings	Celeb_v2	DF40	DiFF		
(1) Img	72.1	74.2	66.0		
(2) SPR	75.7	78.8	73.9		
(3) Img + SPR	84.5	87.2	83.9		
(4) $Img + SPR + Tex/Dir$	87.5	90.9	86.8		
$\overline{(5) \mathbf{T} = \overline{\mathbf{T}}_{bfm}}$	83.3	85.4	79.9		
$(6) \mathbf{T} = \overline{\mathbf{T}}_{bfm} + \mathbf{B}\beta$	84.1	88.2	82.5		
$(7) \mathbf{T} = \mathbf{T}_{msr}$	87.5	90.9	86.8		
(8) w/o shape norm	82.1	83.7	81.0		
(9) w/ shape norm	87.5	90.9	86.8		

Effectiveness of Retinex-based Texture Extraction. Rows 5-7 compare three texture extraction methods (Fig. 3). The Retinex-based T_{msr} achieves optimal performance by capturing fine-grained facial textures, enabling more precise separation of illumination components.

Effectiveness of shape norm. Rows 8-9 demonstrate that shape normalization improves performance by standardizing illumination directions, surface normals, and viewing directions, thereby enhancing the correlation capture between specular reflection, texture, and direct light components under the Phong illumination model.

5 CONCLUSION

In this paper, we build upon the principle that face attributes requiring more estimated parameters and more complex physical constraints are inherently harder to replicate accurately, and demonstrate that the specular reflection component in the Phong illumination model contains generalizable forgery evidence. To enable accurate illumination separation, we develop a faster and more accurate Retinex-based texture extraction method that facilitates precise specular reflection estimation. We then propose SRI-Net to capture the correlations between specular reflection, face texture, and direct light components, thereby enabling the extraction of more complete forgery evidence. Overall, we propose a novel approach for detecting forgery evidence under physically realistic constraints, demonstrating superior performance on both traditional deepfake datasets and more challenging generative deepfake datasets.

REFERENCES

- 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), pp. 1–7. IEEE, 2018.
- Zhongjie Ba, Qingyu Liu, Zhenguang Liu, Shuang Wu, Feng Lin, Li Lu, and Kui Ren. Exposing the deception: Uncovering more forgery clues for deepfake detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 719–728, 2024.
- 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.
- Junyi Cao, Chao Ma, Taiping Yao, Shen Chen, Shouhong Ding, and Xiaokang Yang. End-to-end reconstruction-classification learning for face forgery detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4113–4122, 2022.
- Liang Chen, Yong Zhang, Yibing Song, Lingqiao Liu, and Jue Wang. Self-supervised learning of adversarial example: Towards good generalizations for deepfake detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 18710–18719, 2022.
- Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*, 2025.
- Harry Cheng, Yangyang Guo, Tianyi Wang, Liqiang Nie, and Mohan Kankanhalli. Diffusion facial forgery detection. In *Proceedings of the 32nd ACM international conference on multimedia*, pp. 5939–5948, 2024a.
- Jikang Cheng, Zhiyuan Yan, Ying Zhang, Yuhao Luo, Zhongyuan Wang, and Chen Li. Can we leave deepfake data behind in training deepfake detector? *Advances in Neural Information Processing Systems*, 37:21979–21998, 2024b.
- François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings* of the IEEE conference on computer vision and pattern recognition, pp. 1251–1258, 2017.
- Xinjie Cui, Yuezun Li, Ao Luo, Jiaran Zhou, and Junyu Dong. Forensics adapter: Adapting clip for generalizable face forgery detection. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 19207–19217, 2025.
- 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*, pp. 5781–5790, 2020.
- Shichao Dong, Jin Wang, Renhe Ji, Jiajun Liang, Haoqiang Fan, and Zheng Ge. Implicit identity leakage: The stumbling block to improving deepfake detection generalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 3994–4004, 2023.
- Xiaodan Du, Nicholas Kolkin, Greg Shakhnarovich, and Anand Bhattad. Generative models: What do they know? do they know things? let's find out! *arXiv preprint arXiv:2311.17137*, 2023.
- Akinori F Ebihara, Kazuyuki Sakurai, and Hitoshi Imaoka. Specular-and diffuse-reflection-based face spoofing detection for mobile devices. In 2020 IEEE International Joint Conference on Biometrics (IJCB), pp. 1–10. IEEE, 2020.
- Google AI. Contributing data to deepfakedetection. https://ai.googleblog.com/2019/09/contributing-data-to-deepfakedetection.html, 2019. URL https://ai.googleblog.com/2019/09/contributing-data-to-deepfakedetection.html. Accessed: 2021-11-13.
- Jianzhu Guo, Xiangyu Zhu, Yang Yang, Fan Yang, Zhen Lei, and Stan Z Li. Towards fast, accurate and stable 3d dense face alignment. In *European Conference on Computer Vision*, pp. 152–168. Springer, 2020.

- Alexandros Haliassos, Rodrigo Mira, Stavros Petridis, and Maja Pantic. Leveraging real talking faces via self-supervision for robust forgery detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 14950–14962, 2022.
 - Shu Hu, Yuezun Li, and Siwei Lyu. Exposing gan-generated faces using inconsistent corneal specular highlights. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2500–2504. IEEE, 2021.
 - Yonghyun Jeong, Doyeon Kim, Youngmin Ro, and Jongwon Choi. Frepgan: robust deepfake detection using frequency-level perturbations. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pp. 1060–1068, 2022.
 - Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6007–6017, 2023.
 - Kihong Kim, Yunho Kim, Seokju Cho, Junyoung Seo, Jisu Nam, Kychul Lee, Seungryong Kim, and KwangHee Lee. Diffface: Diffusion-based face swapping with facial guidance. *Pattern Recognition*, 163:111451, 2025.
 - Edwin H Land and John J McCann. Lightness and retinex theory. *Journal of the Optical society of America*, 61(1):1–11, 1971.
 - 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*, pp. 5001–5010, 2020a.
 - Yuezun Li and Siwei Lyu. Exposing deepfake videos by detecting face warping artifacts. *arXiv* preprint arXiv:1811.00656, 2018.
 - Yuezun Li, Xin Yang, Pu Sun, Honggang Qi, and Siwei Lyu. Celeb-df: A large-scale challenging dataset for deepfake forensics. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 3207–3216, 2020b.
 - Honggu Liu, Xiaodan Li, Wenbo Zhou, Yuefeng Chen, Yuan He, Hui Xue, Weiming Zhang, and Nenghai Yu. Spatial-phase shallow learning: rethinking face forgery detection in frequency domain. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 772–781, 2021.
 - Yuchen Luo, Yong Zhang, Junchi Yan, and Wei Liu. Generalizing face forgery detection with high-frequency features. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16317–16326, 2021.
 - Long Ma, Zhiyuan Yan, Yize Chen, Jin Xu, Qinglang Guo, Hu Huang, Yong Liao, and Hui Lin. From specificity to generality: Revisiting generalizable artifacts in detecting face deepfakes. *arXiv* preprint arXiv:2504.04827, 2025.
 - Iacopo Masi, Aditya Killekar, Royston Marian Mascarenhas, Shenoy Pratik Gurudatt, and Wael AbdAlmageed. Two-branch recurrent network for isolating deepfakes in videos. In *European conference on computer vision*, pp. 667–684. Springer, 2020.
 - Nguyen, Huy H, Junichi Yamagishi, and Isao Echizen. Capsule-forensics: Using capsule networks to detect forged images and videos. In *ICASSP 2019-2019 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pp. 2307–2311. IEEE, 2019.
 - Yunsheng Ni, Depu Meng, Changqian Yu, Chengbin Quan, Dongchun Ren, and Youjian Zhao. Core: Consistent representation learning for face forgery detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12–21, 2022.
 - Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3d face model for pose and illumination invariant face recognition. In 2009 sixth IEEE international conference on advanced video and signal based surveillance, pp. 296–301. Ieee, 2009.

- 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.
 - 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*, pp. 86–103. Springer, 2020.
 - Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PmLR, 2021.
 - Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.
 - Felix Rosberg, Eren Erdal Aksoy, Fernando Alonso-Fernandez, and Cristofer Englund. Facedancer: Pose-and occlusion-aware high fidelity face swapping. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 3454–3463, 2023.
 - 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/CVF international conference on computer vision*, pp. 1–11, 2019.
 - Clemens Seibold, Anna Hilsmann, and Peter Eisert. Reflection analysis for face morphing attack detection. In 2018 26th European Signal Processing Conference (EUSIPCO), pp. 1022–1026. IEEE, 2018.
 - Rui Shao, Tianxing Wu, and Ziwei Liu. Detecting and recovering sequential deepfake manipulation. In *European Conference on Computer Vision*, pp. 712–728. Springer, 2022.
 - Kaede Shiohara and Toshihiko Yamasaki. Detecting deepfakes with self-blended images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 18720–18729, 2022.
 - Hui Sun, Cheng Li, Boqiang Liu, Zaiyi Liu, Meiyun Wang, Hairong Zheng, David Dagan Feng, and Shanshan Wang. Aunet: attention-guided dense-upsampling networks for breast mass segmentation in whole mammograms. *Physics in Medicine & Biology*, 65(5):055005, 2020.
 - Ke Sun, Shen Chen, Taiping Yao, Ziyin Zhou, Jiayi Ji, Xiaoshuai Sun, Chia-Wen Lin, and Rongrong Ji. Towards general visual-linguistic face forgery detection. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 19576–19586, 2025.
 - Chuangchuang Tan, Yao Zhao, Shikui Wei, Guanghua Gu, Ping Liu, and Yunchao Wei. Frequency-aware deepfake detection: Improving generalizability through frequency space domain learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 5052–5060, 2024.
 - Chengrui Wang and Weihong Deng. Representative forgery mining for fake face detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 14923–14932, 2021.
 - 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/CVF conference on computer vision and pattern recognition*, pp. 8695–8704, 2020.
 - Zhiyuan Yan, Yong Zhang, Yanbo Fan, and Baoyuan Wu. Ucf: Uncovering common features for generalizable deepfake detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 22412–22423, 2023.
 - Zhiyuan Yan, Yuhao Luo, Siwei Lyu, Qingshan Liu, and Baoyuan Wu. Transcending forgery specificity with latent space augmentation for generalizable deepfake detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8984–8994, 2024a.

- Zhiyuan Yan, Taiping Yao, Shen Chen, Yandan Zhao, Xinghe Fu, Junwei Zhu, Donghao Luo, Chengjie Wang, Shouhong Ding, Yunsheng Wu, et al. Df40: Toward next-generation deepfake detection. *Advances in Neural Information Processing Systems*, 37:29387–29434, 2024b.
- Guanqi Zhan, Chuanxia Zheng, Weidi Xie, and Andrew Zisserman. A general protocol to probe large vision models for 3d physical understanding. *Advances in Neural Information Processing Systems*, 37:43468–43498, 2024.
- 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.
- Xiangyu Zhu, Hao Wang, Hongyan Fei, Zhen Lei, and Stan Z Li. Face forgery detection by 3d decomposition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2929–2939, 2021.

A APPENDIX

A.1 THE USE OF LLM

We used Large Language Models (LLMs) to assist with writing polishing and writing refinement during the preparation of this manuscript. The LLMs were employed solely for improving grammar, sentence structure, and overall readability. All research ideas, data analysis, conclusions, and substantive content remain entirely the original work of us. We take full responsibility for the accuracy and integrity of all content presented in this paper.

A.2 IMPLEMENT DETAILS

For training specifications, we employ Xception as the backbone, optimized with the Adam optimizer at an initial learning rate of 1e-3 for 80 epochs. During training on the FF++ dataset, we uniformly sample 32 frames per video sequence, and the input size is 256x256. Data augmentation strategies include HorizontalFlip, RandomCutout and AddGaussianNoise. The multiscale Retinex-based texture extraction method utilizes σ values of [15, 80, 120] for face texture extraction.

A.3 Examples of Specular Reflection Extraction

In this section, we present additional samples of specular reflection extraction in Fig. 4. We utilize 3DDFA for 3D shape extraction and propose a fast and accurate Retinex-based method for texture extraction. Subsequently, we employ spherical harmonic model to fit ambient light and direct light, extracting specular reflection through a residual-based approach under Retinex-based texture constraints.

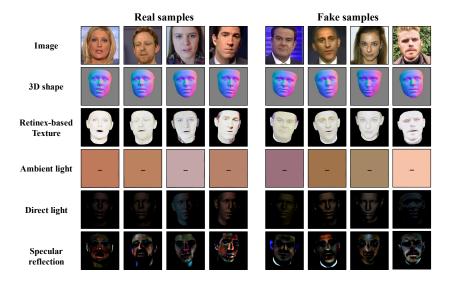


Figure 4: The visualization of Specular Reflection Extraction. The face image can be decomposed into 3D shape, Retinex-based texture, ambient light, direct light, and specular reflection under Phong illumination model constraints. The samples on the left are real samples, while the samples on the right are fake samples.

A.4 LIMITATION ANALYSIS AND FUTURE WORK

Fig. 5 presents cases of misclassification. It can be observed that these cases are characterized by extreme facial poses or significant self-occlusion on the surface, resulting in inaccurate 3D shape fitting and consequently imprecise specular reflection extraction. To address this issue, future work could consider implementing secondary correction of the 3D shape, such as aligning the 3D shape with 2D facial landmarks.

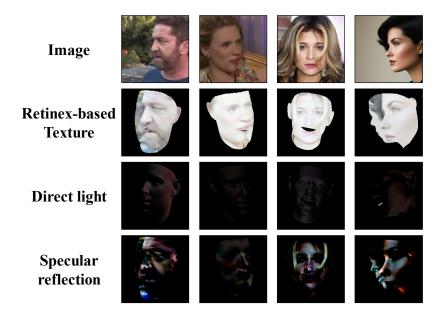


Figure 5: The visualization of misclassified cases. These cases are characterized by extreme facial poses or severe self-occlusion of the facial surface.