

Decoupling Forgery Semantics for Generalizable Deepfake Detection

Wei Ye
weiye@email.ncu.edu.cn

Nanchang University
Nanchang, China

Xinan He
shahur@email.ncu.edu.cn

Feng Ding*
fengding@ncu.edu.cn

Abstract

In this paper, we propose a novel method for detecting DeepFakes, enhancing the generalization of detection through semantic decoupling. There are now multiple DeepFake forgery technologies that not only possess unique forgery semantics but may also share common forgery semantics. The unique forgery semantics and irrelevant content semantics may promote over-fitting and hamper generalization for DeepFake detectors. For our proposed method, after decoupling, the common forgery semantics could be extracted from DeepFakes, and subsequently be employed for developing the generalizability of DeepFake detectors. Also, to pursue additional generalizability, we designed an adaptive high-pass module and a two-stage training strategy to improve the independence of decoupled semantics. Evaluation on FF++, Celeb-DF, DFD, and DFDC datasets shows our method's excellent detection and generalization performance. Code is available at: <https://github.com/leaffeeall/DFS-GDD>.

1 Introduction

In recent years, the emergence of Deepfake technology has introduced significant challenges to the authenticity and trustworthiness of visual content online. These sophisticated algorithms can produce synthetic facial images and videos with remarkable realism, making it increasingly difficult to distinguish between genuine and fake media. Albeit the development of various deepfake detection methods [1, 2, 3, 4, 5, 6, 7, 8], their performance often suffers when applied to real-world scenarios due to a lack of robustness [9] and generalizability [10].

Existing methods typically fall into two categories [11]: direct classification and shallow decoupling classification. However, these approaches often suffer from interference by irrelevant content semantics or overlook common and unique forgery semantics, leading to limited generalization [12], as illustrated in Figure 1. They often perform well in intra-domain detection, i.e., with testing datasets similar to the training dataset, but experience severe performance degradation in cross-domain detection, i.e., with testing datasets different from the training dataset.

*Corresponding author

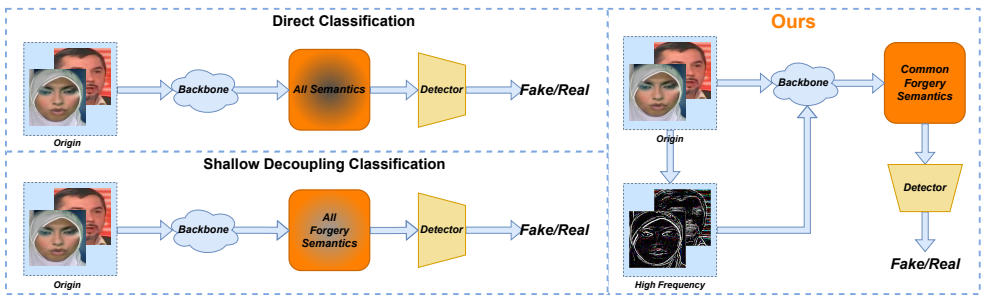


Figure 1: Comparison of our method with Existing Techniques. Orange represents semantics enhancing generalization, while grey indicates hindering generalization.

Also, relying solely on semantics extracted from images for training can lead to detectors overly relying on color textures specific to certain forgery methods [23]. Although some methods attempt to incorporate frequency domain data through techniques like Fourier transforms [5] and Discrete Cosine Transform [64], they often ignore the correlation between frequency domain and traditional color textures.

To address these challenges, we propose a method that leverages high-frequency features and deep decoupling to extract common forgery semantics, unique forgery semantics, and irrelevant content semantics, separating them as independent semantics for forensics purposes. Our approach consists of two training stages. In the first stage, high-frequency features along with images are utilized as inputs. We employ multi-scale high-frequency feature extraction and fusion modules, leveraging vision transformers for global feature extraction. In the second stage, we refine common forgery semantics to enhance generalizable detection performance.

Our contributions are as follows:

1. We propose a forensics model that leverages high-frequency features using our designed adaptive high-pass filter (AHF), combined with deep decoupling to extract common forgery semantics for generalizable DeepFake detection.
2. We introduce two modules: a multi-scale high-frequency feature extraction module (MHFE) and a multi-scale high-frequency feature fusion module (MHFF) to improve the independence and efficiency of extracted forgery semantics.
3. Other than the satisfying intra-domain detection performance, we demonstrate the superior generalization capability of our model in cross-domain deepfake detection scenarios through sufficient evaluations. Also, the impacts of the proposed modules are validated with ablation studies.

The remainder of the paper is organized as follows. In the following section, we survey related works. The proposed method is described in Section 3 and the evaluation results are reported in Section 4. At last, we conclude the paper.

2 Related Work

Deepfake detection. Deepfake detection methods are mainly divided into spatial-based forgery detection and frequency-based forgery detection. Spatial-based forgery detection primarily focuses on examining appearance features in the spatial domain [12, 17, 24]. These

methods often achieve satisfactory performance in intra-domain evaluation but encounter significant performance degradation in cross-domain testing. Some important features are difficult to be discovered when only utilizing RGB/spatial information, but these clues are often better revealed in the frequency domain [40]. Therefore, some methods attempt to improve generalization by leveraging frequency domain components [4, 6, 23, 25, 56]. However, the improvement of these techniques in generalizing to unknown forgery technologies remains limited [2, 8, 10, 12].

Decoupling in Deepfake Detection. Decoupling, a technique breaking down complex semantics into simpler, more discriminative variables, has gained attention [21]. Some studies have attempted to decouple forgery semantics for detection [8, 10, 16]. Additionally, Liang *et al.* [21] reinforced feature independence through content consistency and global representation contrastive constraints. Recent research [22, 35] further attempts to decompose manipulation-related semantics into unique and common forgery semantics, utilizing common forgery semantics for detection, alleviating the issue of performance degradation in cross-domain detection.

3 Proposed Method

3.1 Motivation

We aim to improve the generalization performance of deepfake detection by addressing three main challenges. Firstly, irrelevant content semantics may lead to the over-fitting of detectors, thereby hindering generalization. Secondly, the diversity of forgery techniques leads to distinct forgery artifacts, making detectors trained on specific sets of artifacts less effective at identifying unseen forgeries. Additionally, relying solely on RGB information may lead to over-reliance on specific color textures, while extracting features in a multi-modal manner can make the method more effective [23, 30, 56].

3.2 Overview

Our method involves two-stage training, outlined in Figure 2. In training stage 1, a pair of real and fake images along with their corresponding high-frequency features are inputs. *Encoder*₁ decouples them into irrelevant content semantics and all forgery semantics *Fa*, which are then identified by *Detector*₁. This stage ensures feature disentanglement through self-reconstruction and cross-reconstruction using *Decoder*₁. *Fa* is composed of common forgery semantics *Fc* and unique forgery semantics *Fu*, intertwined as follows:

$$Fa = [Fc, Fu] \quad (1)$$

In training stage 2, the identical real and fake image pairs are employed as inputs again. *Encoder*₂ extracts *Fa* from *Encoder*₁, further disentangling them into *Fc* and *Fu*. *Detector*₂ identifies *Fu*, while *Detector*₃ identifies *Fc*. This stage ensures semantics disentanglement through self-reconstruction and cross-reconstruction using *Decoder*₂. During inference, only *Fc* from *Encoder*₂ are used for real-fake classification, followed by *Detector*₃.

3.3 Training Stage 1

Encoder1. SwiftFormer [29] introduces an efficient additive attention mechanism to replace costly matrix operations, achieving a better balance between accuracy and efficiency. We

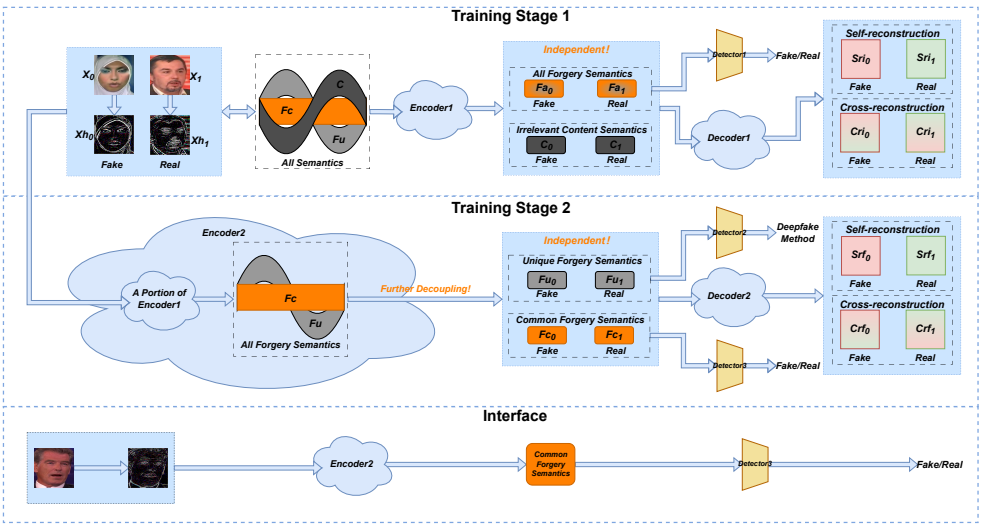


Figure 2: The overview of our method. In the entangled semantics, dark gray represents irrelevant content semantics C , light gray represents unique forgery semantics F_u , and orange represents common forgery semantics F_c . $Encoder_2$ utilizes branches to extract all forgery semantics from $Encoder_1$. Both $Decoder_1$ and $Decoder_2$ include processes for self-reconstruction and cross-reconstruction.

use SwiftFormer-L1 to extract irrelevant content from images. Additionally, we employ Xception [6] for RGB and high-frequency feature extraction. To improve the efficiency of high-frequency features, we design an adaptive high-pass filter (AHF), inspired by the Gaussian filter. The convolution kernel of AHF is devised using the following formula, overcoming the fixed parameter limitations in traditional high-pass filtering:

$$g(x, y, \sigma) = E - \frac{G(x, y, \sigma)}{\sum_{x, y} G(x, y, \sigma)} \quad (2)$$

where g represents the convolution kernel of AHF, x and y denote positions in the convolution kernel, σ is set to 1.0, E is a matrix with a center value of 1 and the rest as 0, G is the Gaussian distribution function.

Furthermore, to ensure the high-pass filtering characteristics of AHF, after initialization and each backpropagation, we reset the center element to -1 and normalize the remaining elements:

$$\begin{cases} \hat{g}(x, y) = -1 & \text{if } (x, y) = (0, 0) \\ \hat{g}(x, y) = \frac{g(x, y)}{\sum_{x, y} g(x, y) - g(0, 0)} & \text{if } (x, y) \neq (0, 0) \end{cases} \quad (3)$$

where \hat{g} represents the updated convolution kernel of AHF, g represents the original convolution kernel of AHF, and x and y denote positions in the convolution kernel, with $(x, y) = (0, 0)$ indicating the center of the kernel.

To enhance the effectiveness of semantics for forensics purposes, we design a multi-scale high-frequency feature extraction module (MHFE) based on the adaptive high-pass filter. We also employ the Pag feature fusion method [64] to design a multi-scale high-frequency feature fusion module (MHFF), which allows us to integrate high-frequency features into the

RGB stream, preventing over-reliance on specific color textures. In Section 4, we conducted ablation studies to justify the effectiveness of the proposed modules. The supplementary material illustrates the network architecture of $Encoder_1$ and demonstrates the process for inputs X_1 and Xh_1 . The entire process is represented as:

$$Encoder_1(X_i, Xh_i) = C_i, Fa_i \quad (4)$$

where $i = (0, 1)$ denotes the image index. X_i , Xh_i , C_i , and Fa_i represent the RGB information, high-frequency features, irrelevant content semantics, and all forgery semantics, respectively.

Decoder1. We designed a dual-channel network in $Decoder_1$. One channel simultaneously utilizes convolutional layers and SwiftFormer [24] to process irrelevant content semantics, while the other channel solely adopts convolutional layers to handle common and unique forgery semantics. These semantics are fused to reconstruct the image, enhancing decoupling capability using self-reconstruction and cross-reconstruction techniques [33]. The supplementary material illustrates the network framework of $Decoder_1$ and demonstrates the process for inputs C_1 and Fa_1 . The overall process is represented as:

$$Decoder_1(C_{i_1}, Fa_{i_2}) = \begin{cases} Sri_{i_1} & \text{if } i_1 = i_2 \\ Cri_{i_1} & \text{if } i_1 \neq i_2 \end{cases} \quad (5)$$

where $i_1 = (0, 1)$ and $i_2 = (0, 1)$ denote the image indices. C_{i_1} , Fa_{i_2} , Sri_{i_1} , and Cri_{i_1} represent irrelevant content semantics, all forgery semantics, self-reconstructed images, and cross-reconstructed images, respectively.

Objective Function. The framework’s overall loss function combines two distinct components using a weighted sum: forgery semantics detection loss and reconstruction loss.

Classification Loss. For detecting all forgery semantics, we use cross-entropy loss:

$$L_{cls} = L_{ce}(Detector_1(Fa_i), y_i) \quad (6)$$

where $i = (0, 1)$ denotes the image indices. L_{ce} , Fa_i , and y_i (fake, real) represent the cross-entropy loss, all forgery semantics, and binary classification labels, respectively.

Reconstruction Loss. We employ an L_1 reconstruction loss to maintain feature completeness and image consistency:

$$L_{rec} = \|X_{i_1} - Decoder_1(C_{i_1}, Fa_{i_2})\|_1 \quad (7)$$

where $i_1 = (0, 1)$ and $i_2 = (0, 1)$ denote the image indices. X_{i_1} , C_{i_1} , and Fa_{i_2} represent the original image, irrelevant content semantics, and all forgery semantics, respectively.

Overall Loss. The final loss function L_{s_1} for stage 1 is:

$$L_{s_1} = \rho_1 L_{cls} + \rho_2 L_{rec} \quad (8)$$

where ρ_1 and ρ_2 are trade-off hyperparameters.

3.4 Training Stage 2

Encoder2. In $Encoder_2$, we utilize a portion of $Encoder_1$ responsible for extracting all forgery semantics, followed by convolutional layers to further disentangle these semantics

into unique and common forgery semantics. The supplementary material illustrates the process for inputs X_1 and Xh_1 . The overall process is represented as:

$$Encoder_2(X_i, Xh_i) = Fu_i, Fc_i \quad (9)$$

where $i = (0, 1)$ denotes the image indices. X_i , Xh_i , Fu_i , and Fc_i represent RGB information, high-frequency features, unique forgery semantics, and common forgery semantics, respectively.

Decoder2. In *Decoder₂*, we recognize that unique and common forgery semantics contain more local information. Hence, we design two dual-channel networks, each comprising only convolutional layers, and merge them during the process to reconstruct image semantics. We continue to use self-reconstruction and cross-reconstruction to improve decoupling capability [53]. The supplementary material illustrates the process for inputs Fu_1 and Fc_1 to obtain self-reconstructed image semantics Srf_1 . The entire process is summarized as:

$$Decoder_2(Fu_{i_1}, Fc_{i_2}) = \begin{cases} Srf_{i_1} & \text{if } i_1 = i_2 \\ Crf_{i_1} & \text{if } i_1 \neq i_2 \end{cases} \quad (10)$$

where $i_1 = (0, 1)$ and $i_2 = (0, 1)$ denote the image indices. Fu_{i_1} , Fc_{i_2} , Srf_{i_1} , and Crf_{i_1} represent unique forgery semantics, common forgery semantics, self-reconstructed image semantics, and cross-reconstructed image semantics, respectively.

Objective Function. The overall loss function of the framework combines four distinct components using a weighted sum: unique forgery semantics detection loss, common forgery semantics detection loss, contrastive loss, and reconstruction loss.

Classification Loss. For detecting unique forgery semantics, we use cross-entropy loss:

$$L_{cls_1} = L_{ce}(Detector_2(Fu_i), S_i) \quad (11)$$

where $i = (0, 1)$ denotes the image indices. L_{ce} , Fu_i , and S_i represent the cross-entropy loss, unique forgery semantics, and multi-class label regarding the forgery method, respectively.

For detecting common forgery semantics, we also use cross-entropy loss:

$$L_{cls_2} = L_{ce}(Detector_3(Fc_i), y_i) \quad (12)$$

where $i = (0, 1)$ denotes the image indices. L_{ce} , Fc_i , and y_i (fake, real) represent the cross-entropy loss, common forgery semantics, and binary classification labels, respectively.

The total classification loss is:

$$L_{cls} = \rho_3 L_{cls_1} + \rho_4 L_{cls_2} \quad (13)$$

where ρ_3 and ρ_4 are trade-off hyperparameters.

Contrastive Loss. To enhance the encoding capability of the encoder for different image semantics, we utilize contrastive loss L_{con} . The loss function minimizes the distance between the anchor image semantics f_i^a and its corresponding positive image semantics f_i^+ , while simultaneously maximizing the distance between the anchor image semantics and its corresponding negative image semantics f_i^- [53]. Here, positive samples indicate the same source, and negative samples indicate different sources. For real images, the same source refers to other real images, and different sources refer to fake images. For fake images, the

same source refers to other fake images created using the same forgery method, and different sources refer to real images. The formula is as follows:

$$L_{con} = \max\{0, a + \|f_i^a - f_i^+\|_2 - \|f_i^a - f_i^-\|_2\} \quad (14)$$

where $i = (0u, 0c, 1u, 1c)$ represents unique forgery semantics for fake images, common forgery semantics for fake images, unique forgery semantics for real images, and common forgery semantics for real images, respectively. a is a trade-off hyperparameter.

Reconstruction Loss. Self-reconstruction and cross-reconstruction are employed to reconstruct image semantics with reconstruction loss L_{rec} that can be written as :

$$L_{rec} = \|Fa_{i_1} - Decoder_2(Fc_{i_1}, Fu_{i_2})\|_1. \quad (15)$$

where $i_1 = (0, 1)$ and $i_2 = (0, 1)$ denote image indices. Fa_{i_1} , Fc_{i_1} , and Fu_{i_2} represent all forgery semantics, common forgery semantics, and unique forgery semantics, respectively.

Overall Loss. The final loss function L_{s_2} for the training stage 2 is:

$$L_{s_2} = L_{cls} + \rho_5 L_{con} + \rho_6 L_{rec} \quad (16)$$

where ρ_5 and ρ_6 are trade-off hyperparameters.

4 Experimental Settings

Datasets. We trained our model on FaceForensics++ (FF++) [27] and evaluated it on FF++, DeepfakeDetection (DFD) [13], Deepfake Detection Challenge (DFDC) [10], and Celeb-DF [20]. FF++ includes images generated by five facial manipulation algorithms: DeepFakes (DF) [8], Face2Face (F2F) [32], FaceSwap (FS) [18], NeuralTexture (NT) [33], and FaceShifter (FST) [19].

Implementation. We utilized PyTorch and NVIDIA RTX 3090Ti for training. Images were resized to 256×256 , while all models were trained for 20 epochs with a fixed batch size of 16. We employed SGD optimizer [26] with a learning rate of $\beta = 5 \times 10^{-4}$. In training stage 1, we set $\rho_1 = 1.0$ and $\rho_2 = 0.3$ in Equation (8). For training stage 2, we set $\rho_3 = 0.1$ and $\rho_4 = 1.0$ in Equation (13), $a = 3.0$ in Equation (14), and $\rho_5 = 0.05$ and $\rho_6 = 0.3$ in Equation (16).

Evaluation Metrics. We employ the Area Under Curve (AUC) metric for performance evaluation, consistent with prior studies [6, 15, 22, 23, 25, 31, 35].

4.1 Results

Intra-domain Sub-datasets Performance: Our approach, as depicted in Table 1, effectively separates domain-specific forgery, mitigating over-fitting risks. Compared to Resnet-50 [15], EfficientNet-B4 [30], Xception [6], SRM [23], F3-Net [25], UCF [35], and Lin *et al.* [22], our method consistently achieves higher AUC scores across all sub-datasets.

Cross-domain Datasets Performance: Trained on FF++, our method outperform other benchmarks when tested on FF++ [32], Celeb-DF [18], DFD [33], and DFDC [10], as summarized in Table 2. This demonstrates superior generalization capabilities, resulting in the most accurate detection outcomes.

Method	AUC(%)				
	F2F [52]	FS [48]	NT [33]	DF [0]	FST [19]
ResNet-50 [15]	93.76	93.30	83.43	93.34	92.25
EfficientNet-B4 [31]	97.41	97.10	90.87	97.02	96.28
Xception [9]	96.92	95.85	94.00	97.47	95.62
SRM [23]	96.49	97.59	92.66	97.64	97.55
F3-Net [25]	96.56	94.14	93.15	97.67	96.80
UCF [35]	97.12	97.46	91.99	97.40	97.31
Lin <i>et al.</i> [22]	98.37	97.97	95.06	98.86	98.41
Ours	99.15	99.36	96.23	99.29	99.13

Table 1: Methods trained on FF++ are evaluated intra-domain on sub-datasets separated by five types of forgeries: F2F, FS, NT, DF, and FST. Best results are highlighted in bold.

Method	AUC(%)			
	FF++ [32]	Celeb-DF [48]	DFD [33]	DFDC [0]
ResNet-50 [15]	91.06	64.78	72.94	53.38
EfficientNet-B4 [31]	95.63	67.80	76.81	56.59
Xception [9]	95.93	69.37	78.08	56.87
SRM [23]	96.30	68.08	77.57	58.22
F3-Net [25]	95.64	67.62	80.51	55.96
UCF [35]	96.17	70.48	75.68	55.20
Lin <i>et al.</i> [22]	97.68	75.19	80.56	62.18
Ours	98.58	76.94	83.02	62.55

Table 2: Methods trained on FF++ and tested on multiple datasets, including Celeb-DF, DFD, and DFDC, are presented. Best results are highlighted in bold.

4.2 Ablation Study

We evaluated the model combining RGB and high-frequency features along with its variations of multi-scale high-frequency feature extraction (MHFE) and multi-scale high-frequency feature fusion (MHFF) on the FF++ database. RGB represents using only RGB images, and High-frequency represents using only high-frequency features. All models were trained on FF++ and tested on four datasets, as shown in Table 3, demonstrating the complementary nature of the two modalities and the effectiveness of each module.

4.3 Visualization

For a more intuitive demonstration of our method’s effectiveness, we visualize the Grad-CAM [24] of Xception, UCF, and our method, as shown in Figure 3. Grad-CAM reveals that Xception tends to overfit to small local regions or focus on content noises outside facial regions when unconstrained. Although UCF performs well on certain irrelevant facial content semantics, it fails in certain cases, particularly when irrelevant facial semantics represent individuals with a dark skin tone, hindering its generalization performance. In contrast, our method consistently focuses on tracing the common clues in the DeepFakes left by different generating models, regardless of irrelevant facial content semantics, showcasing outstanding

Method	AUC(%)			
	FF++ [62]	Celeb-DF [18]	DFD [63]	DFDC [9]
RGB	97.77	74.80	79.20	61.18
High-frequency	96.90	73.55	81.57	60.75
RGB + High-frequency(Fusion)	98.33	73.70	82.19	61.11
Fusion + MHFE	98.50	74.72	80.03	60.92
Fusion + MHFF	98.02	72.72	80.44	61.18
Fusion + MHFE + MHFF	98.58	76.94	83.02	62.55

Table 3: Ablation study on FF++. The best results are highlighted in bold.

generalization performance.

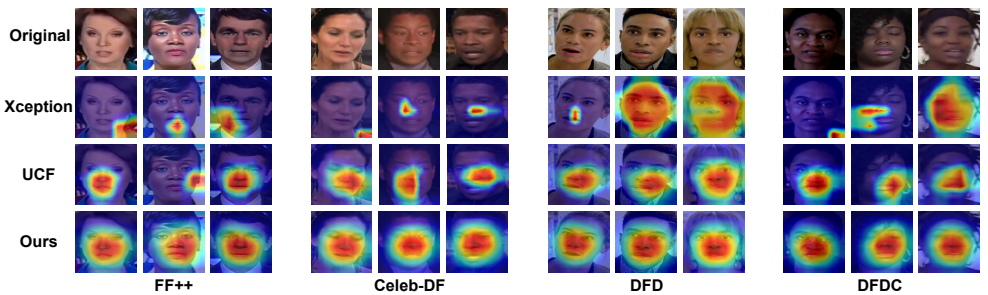


Figure 3: Visualization of Grad-CAM for Xception, UCF, and our approach across intra-domain (FF++) and cross-domain datasets (Celeb-DF, DFD, and DFDC).

5 Conclusion

In this paper, we propose a semantics decoupling approach for training a DeepFake detector, achieving satisfying forensics performance. In particular, by extracting and analyzing common forgery semantics of different DeepFake technologies, the proposed method is validated to be capable of highly generalizable DeepFake detection. Also, justified by the ablation study, the modules designed in our proposed method effectively enhance both the intra-domain and cross-domain detection performance. Our work also provides insights for discerning AI-generated content by employing semantics disentanglement. Enhancing the independence and effectiveness of forgery semantics may be crucial for obtaining high-quality forensics models.

Acknowledgement

This work was supported in part by the National Natural Science Foundation of China under Grant 62262041, and in part by the Jiangxi Provincial Natural Science Foundation under Grant 20232BAB202011.

References

- [1] Deepfake detection challenge. <https://www.kaggle.com/c/deepfake-detection-challenge>. Accessed: 2021-04-24.
- [2] Deepfakes. In <https://github.com/deepfakes/faceswap>, 2017.
- [3] 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*, pages 719–728, 2024.
- [4] Federico Baldassarre, Quentin Debard, Gonzalo Fiz Pontiveros, and Tri Kurniawan Wijaya. Quantitative metrics for evaluating explanations of video deepfake detectors. In *British Machine Vision Conference (BMVC)*, 2022.
- [5] Zehao Chen and Hua Yang. Attentive semantic exploring for manipulated face detection. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1985–1989, 2021.
- [6] 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.
- [7] Feng Ding, Guopu Zhu, Yingcan Li, Xinpeng Zhang, Pradeep K Atrey, and Siwei Lyu. Anti-forensics for face swapping videos via adversarial training. *IEEE Transactions on Multimedia*, 24:3429–3441, 2021.
- [8] Feng Ding, Bing Fan, Zhangyi Shen, Keping Yu, Gautam Srivastava, Kapal Dev, and Shaohua Wan. Securing facial bioinformation by eliminating adversarial perturbations. *IEEE Transactions on Industrial Informatics*, 2022.
- [9] Feng Ding, Zhangyi Shen, Guopu Zhu, Sam Kwong, Yicong Zhou, and Siwei Lyu. Exsgan: Synthesizing anti-forensics images via extra supervised gan. *IEEE Transactions on Cybernetics*, 2022.
- [10] Bing Fan, Shu Hu, and Feng Ding. Synthesizing black-box anti-forensics deepfakes with high visual quality. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4545–4549, 2024.
- [11] Zhixiao Fu, Xinyuan Chen, Daizong Liu, Xiaoye Qu, Jianfeng Dong, Xuhong Zhang, and Shouling Ji. Multi-level feature disentanglement network for cross-dataset face forgery detection. *Image and Vision Computing*, 135:104686, 2023.
- [12] Yueying Gao, Weiguo Lin, Junfeng Xu, Wanshan Xu, and Peibin Chen. Self-supervised adversarial training for robust face forgery detection. In *British Machine Vision Conference (BMVC)*, 2023.
- [13] Google and Jigsaw. Deepfakes dataset by google & jigsaw. In <https://ai.googleblog.com/2019/09/contributing-data-to-deepfakedetection.html>, 2019.

- [14] Hui Guo, Shu Hu, Xin Wang, Ming-Ching Chang, and Siwei Lyu. Robust attentive deep neural network for detecting gan-generated faces. *IEEE Access*, 10:32574–32583, 2022.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [16] Jiashang Hu, Shilin Wang, and Xiaoyong Li. Improving the generalization ability of deepfake detection via disentangled representation learning. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 3577–3581, 2021.
- [17] Yao-Hui Huang, Jun-Wei Hsieh, Ming-Ching Chang, Lipeng Ke, Siwei Lyu, and Arpita Samanta Santra. Multi-teacher single-student visual transformer with multi-level attention for face spoofing detection. In *British Machine Vision Conference (BMVC)*, 2021.
- [18] Kowalski. Faceswap. In <https://github.com/MarekKowalski/FaceSwap/>, 2018.
- [19] Lingzhi Li, Jianmin Bao, Hao Yang, Dong Chen, and Fang Wen. Faceshifter: Towards high fidelity and occlusion aware face swapping. *arXiv preprint arXiv:1912.13457*, 2019.
- [20] 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*, pages 3207–3216, 2020.
- [21] Jiahao Liang, Huafeng Shi, and Weihong Deng. Exploring disentangled content information for face forgery detection. In *European Conference on Computer Vision*, pages 128–145, 2022.
- [22] Li Lin, Xinan He, Yan Ju, Xin Wang, Feng Ding, and Shu Hu. Preserving fairness generalization in deepfake detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [23] 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*, pages 16317–16326, 2021.
- [24] Wenbo Pu, Jing Hu, Xin Wang, Yuezun Li, Shu Hu, Bin Zhu, Rui Song, Qi Song, Xi Wu, and Siwei Lyu. Learning a deep dual-level network for robust deepfake detection. *Pattern Recognition*, 130:108832, 2022.
- [25] 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, 2020.
- [26] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.

- [27] 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*, pages 1–11, 2019.
- [28] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [29] Abdelrahman Shaker, Muhammad Maaz, Hanoona Rasheed, Salman Khan, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Swiftformer: Efficient additive attention for transformer-based real-time mobile vision applications. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17425–17436, 2023.
- [30] Luchuan Song, Zheng Fang, Xiaodan Li, Xiaoyi Dong, Zhenchao Jin, Yuefeng Chen, and Siwei Lyu. Adaptive face forgery detection in cross domain. In *European Conference on Computer Vision*, pages 467–484, 2022.
- [31] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114, 2019.
- [32] 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.
- [33] 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.
- [34] Jiacong Xu, Zixiang Xiong, and Shankar P. Bhattacharyya. Pidnet: A real-time semantic segmentation network inspired from pid controller. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023.
- [35] 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*, pages 22412–22423, 2023.
- [36] Yue Zhou, Bing Fan, Pradeep K. Atrey, and Feng Ding. Exposing deepfakes using dual-channel network with multi-axis attention and frequency analysis. In *Proceedings of the 2023 ACM Workshop on Information Hiding and Multimedia Security*, pages 169–174, 2023.