

Exploiting Style Latent Flows for Generalizing Deepfake Video Detection

Jongwook Choi^{1,2} Taehoon Kim¹ Yonghyun Jeong³ Seungryul Baek⁴ Jongwon Choi^{1,2*}

¹Dept. of Advanced Imaging, Chung-Ang Univ, Korea ²GS. of AI, Chung-Ang Univ, Korea

³Image Vision, NAVER Cloud, Korea ⁴AI Graduate School, UNIST, Korea

{cjw, kimth}@vilab.cau.ac.kr, yonghyun.jeong@navercorp.com, srbaek@unist.ac.kr, choijw@cau.ac.kr

Abstract

This paper presents a new approach for the detection of fake videos, based on the analysis of style latent vectors and their abnormal behavior in temporal changes in the generated videos. We discovered that the generated facial videos suffer from the temporal distinctiveness in the temporal changes of style latent vectors, which are inevitable during the generation of temporally stable videos with various facial expressions and geometric transformations. Our framework utilizes the StyleGRU module, trained by contrastive learning, to represent the dynamic properties of style latent vectors. Additionally, we introduce a style attention module that integrates StyleGRU-generated features with content-based features, enabling the detection of visual and temporal artifacts. We demonstrate our approach across various benchmark scenarios in deepfake detection, showing its superiority in cross-dataset and cross-manipulation scenarios. Through further analysis, we also validate the importance of using temporal changes of style latent vectors to improve the generality of deepfake video detection.

1. Introduction

Recent generative algorithms are capable of producing high-quality videos; while this advancement has led to social concerns, as it becomes increasingly challenging to distinguish between generated videos and authentic ones. Generative models have the potential to expedite industries such as entertainment, gaming, fashion, design, and education. However, their misuse can have adverse effects on society. The high quality of the videos intensifies the potential for inappropriate utilization of the technique. To resolve the issue, researchers are actively engaged in developing the deepfake video detection algorithms [19, 29, 40].

Early deepfake detection research addressed spatial artifacts such as unnatural aspects [30] and frequency-level checkerboard of the generative model [33] in the genera-

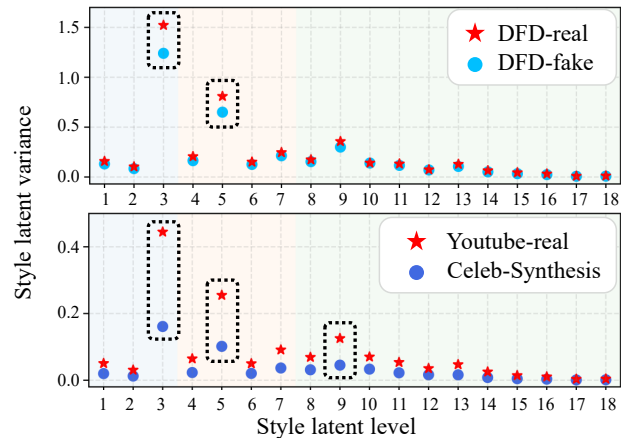


Figure 1. **Variance of style flow for each style latent level.** The x-axis shows the level of style latent vectors for fine style representations. We noticed that the level-wise differences vary across deepfake domains, but the variance of style latent vectors is particularly lower in certain levels of the style latent vectors for fake videos than in real videos. This happens due to the temporal smoothness of the style latent vectors to create temporally stable deepfake videos, and our results demonstrate that deepfake videos have a distinct variance in style flow compared to real videos.

tion of a single frame. While spatial artifact-based detection methods have shown reasonable performance for single images, they failed in accounting for temporal artifacts in deepfake videos that consist of multiple frames. To address the limitation, recent studies [15, 41] integrate temporal cues such as flickering [14, 57] and discontinuity [16] to enhance the accuracy of deepfake video detection task.

Nevertheless, existing methods [28, 51] that exploit visual and temporal cues of generated videos encountered performance degradation when attempting to identify videos produced by recent deepfake generation algorithms, which effectively suppress visual and temporal artifacts. Recent observations [36, 39] have presented a decline in the effectiveness of visual artifact-based algorithms against newly developed high-quality generation techniques [11, 37]. Temporal artifact-based algorithms [16] also experi-

*Corresponding author

enced performance degradation due to the advanced quality of video generation methods [41, 57].

In this paper, we propose to focus on the suppressed variance in the temporal flow of style latent features which are extracted from generated videos, rather than conventional visual and temporal artifacts. The style latent feature encodes the facial appearance and facial movements, such as expressions and geometric transformations, to control the synthesis of facial images. Fig. 1 shows the clear differences in the flow of style latent vectors between the generated and synthetic videos. We noticed that the suppressed variance of the style latent vector occurs due to the temporal smoothness of facial appearance and movements in the generated videos. We believe that the flow of style latent vectors has the potential to be a distinctive feature for generalizing deepfake video detection across various generative models.

In order to encode the dynamics of style latent vectors, we introduce the StyleGRU module in our video deepfake detection framework. This module can effectively capture the temporal variations within style latent vectors, consequently enabling the extraction of robust style-based temporal features. To increase the representation power of style-based temporal features, we further involve the supervised contrastive learning [23] technique in StyleGRU module. We develop the style attention module which effectively integrates the style-based temporal features with the content features that encode the visual and temporal artifacts.

We perform various experiments to evaluate the proposed algorithm for multiple scenarios such as cross-dataset and cross-manipulation scenarios. From the experimental results, our method notably surpasses the existing algorithms, achieving a state-of-the-art performance. Ablation studies are also conducted to show the module-wise effectiveness, and the meaningful analysis experiments confirm that the variations in style latent vector contribute effectively to the robust detection to the various deepfake generation methods. The contributions of our work are summarized as follows.

- We propose a novel video deepfake detection framework that is based on the unnatural variation of the style latent vectors.
- We introduce a StyleGRU module that can encode the variations of style latent vectors with the contrastive learning technique.
- We propose a style attention module that effectively integrates the features from StyleGRU with the content features capturing conventional visual and temporal artifacts.
- Our approach demonstrates state-of-the-art performance in various deepfake detection scenarios, including cross-dataset and cross-manipulation settings, which confirms the effectiveness of facial attribute changes for deepfake video detection.

2. Related work

Deepfake detection. In deepfake detection research, there are two main types: image-level and video-level detectors. The image-level detectors detect fake images by recognizing spatial artifacts of a single frame. A straightforward approach to detect spatial artifacts involves using Xception model [40], a Convolutional Neural Network (CNN) architecture along with MAT [56] that utilizes an attentional mechanism. Face X-ray [27] proposed a method that uses boundaries between forged faces and backgrounds to capture spatial inconsistency. Several methods utilized the frequency-level features to emphasize the artifacts from the image generators [14]. FrePGAN [18] utilized the frequency map transformed from the original RGB image, and LRL [6] integrated RGB images with frequency-aware cues to learn inconsistencies. Recently, algorithms utilize methods including blending artifacts [2, 5, 12, 43] or separating elements relevant to detection from irrelevant ones [12, 55] during their training. Image-level detectors demonstrate high generalization performance. However, image-based methods have the limitation of not utilizing the temporal cues that inevitably occur during the generation of deepfake videos.

In contrast to image-level detectors, video-level detectors take advantage of temporal information by using multiple frames to detect deepfake videos. CNN-GRU [42] captured temporal information from the features extracted by CNNs, and LipForensics [16] utilized the information near the mouth region by employing a network pre-trained on a LipReading dataset. Recently, FTCN [57] directly extracted temporal information using 3D CNNs with a spatial kernel size of 1. AltFreeze [52] showcases strong generalization capabilities through its unique approach of independently training spatial and temporal information. However, the existing methods have focused on low-level features like pixels, so they ignore the temporal changes of high-level contexts such as facial attributes. To improve detection capabilities, we propose a novel approach that considers temporal features at both low-level and high-level, thereby enhancing the overall context and advancing the state-of-the-art in deepfake detection.

Style Latent Vectors. Well-trained Generative Adversarial Network (GAN) models [20–22] successfully generate disentangled latent space. This advancement permits the straightforward manipulation of style latent features to produce desired facial images. The GAN inversion, initially introduced by [59], targets the identification of a latent code within the latent space, which makes it possible to generate an image closely resembling a target one. Furthermore, various encoders [1, 38, 48, 58] have been studied to properly utilize the prior knowledge of StyleGAN latent space. Moreover, GAN inversion methods decoding the style latent features of StyleGAN [21] are utilized for generating the

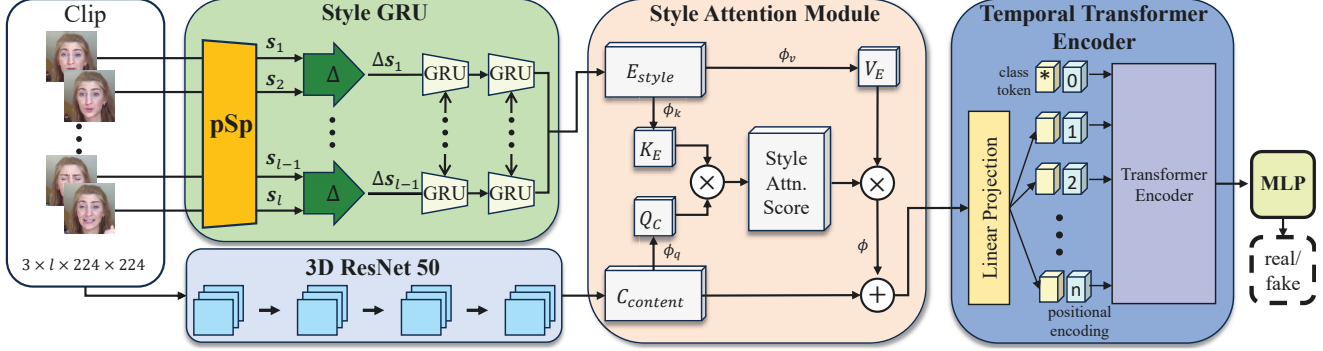


Figure 2. **Schematic diagram of our entire framework.** For the video deepfake detection, we first extract the style latent vectors from pSp for each individual frame, and then encode their variations using the StyleGRU module into the style-based temporal feature E_{style} . In parallel, the content feature $C_{content}$ is extracted via the 3D ResNet-50 architecture from the video clip. Style Attention Module (SAM) integrates the style-based temporal feature E_{style} with the content feature $C_{content}$ via the attention mechanism. Finally, Temporal Transformer Encoder (TTE) is applied to map the representation into the binary class of real and fake labels.

high-quality face-swap videos. MegaFS [60] also inverted both source and target images into the finely-tuned StyleGAN latent space for the purpose of face swap.

Face swapping through GAN inversion offers the advantage of separating facial identity in the latent space and swapping latent features, thereby preserving fidelity in models like StyleGAN [20]. Furthermore, it is employed for various Deepfake generation techniques, including modifying facial attributes, extending beyond just altering facial identity. These approaches surpassed the existing method, such as StyleRig [45], which involved direct modifications to the style latent feature. The GAN inversion-based face swap methods [31, 53, 54] continue to demonstrate good performance. We also employ the style latent features of GAN inversion models to represent the temporal changes of facial attributes. Given that most of the recent GAN inversion models, including those based on the StyleGAN architecture [20], utilize style features, a similar situation applies to the construction of deepfake generation models. Hence, we propose the use of StyleGRU, which is well suited for countering attacks on state-of-the-art deepfake detection models.

3. Methodology

In this section, we propose the video deepfake detection framework that exploits the style latent vectors and captures their characteristics in the generated videos. Especially, we proposed four distinct modules to achieve the goal: (1) StyleGRU module that extracts the pSp features for each individual frame and encode the temporal variations of the style latent vectors using the GRU into the style-based temporal feature E_{style} , (2) 3D ResNet-50 module that encodes the video clip into the content feature $C_{content}$, (3) Style Attention Module (SAM) that combines the style-based temporal feature E_{style} with the content feature $C_{content}$ via the

attention mechanism, (4) Temporal Transformer Encoder (TTE) module that effectively maps the combined features into the binary classes (i.e. real and fake). Fig. 2 illustrates the overall flow of our framework.

3.1. Overall Architecture

3.1.1 Input configuration

We obtain the video clip I that contains a sequence of l consecutive frames extracted from a single video. Then, we build a clip I_p composed of all cropped faces within the input clip I that have been preprocessed with FFHQ [20] conditions. The preprocessed video clip I_p is used as an input to StyleGRU after resizing it to be $\mathcal{R}^{256 \times 256 \times l}$, while the raw clip I is an input of 3D ResNet-50 module with the resized shape of $\mathcal{R}^{224 \times 224 \times l}$.

3.1.2 StyleGRU module

The preprocessed video clip I_p serves as the input to extract style latent vectors via the pSp encoder [38]. The pSp encoder (pSp) provides the capability to extract style latent vectors for inverting given images within a well-trained StyleGAN [21] latent space. We extract the style latent vectors \mathbf{S} as:

$$\mathbf{S} = pSp(I_p) \equiv \{s_1, s_2, \dots, s_l\}, \quad (1)$$

where $s_k \in \mathcal{R}^D$ is the D -dimensional style vector of the k -th frame of the clip I_p , $k \in [1, l]$, and l is the number of frames in clip.

The style latent vector that we extract through the pSp encoder is a feature for the GAN inversion task, so it is insufficient for direct application to the deepfake detection task. To enable the extension of style from the image domain to the video, we employ the difference of subsequent

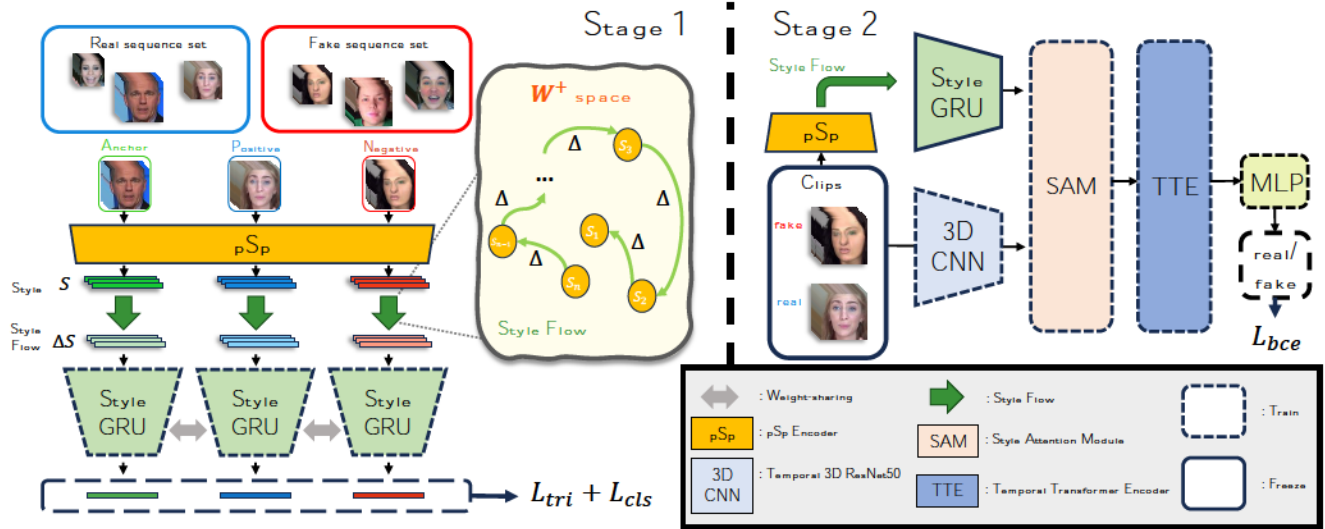


Figure 3. **Our training procedures.** In stage 1, we train the StyleGRU module using the supervised contrastive learning technique to effectively capture variations of the pSp feature and encode it into the robust style-based temporal feature. In stage 2, we train Style Attention Module (SAM) and Temporal Transformer Encoder (TTE) to integrate the style-based temporal feature and content feature and then map it towards the binary classes (ie. real and fake), respectively.

style latent vectors, extracted from frames within a clip. The style flow Δs , derived from difference of style latent vectors, is represented as $\Delta s_i = s_{i+1} - s_i$. Then, the input ΔS , which is called *style flow*, is constructed as follows:

$$\Delta S = \{\Delta s_1, \Delta s_2, \dots, \Delta s_{(l-1)}\}, \quad (2)$$

The style flow ΔS is an input to the Gated Recurrent Unit (GRU) layers [7]. Then, we obtain a temporal embedded style flow $E_{\text{style}} \in \mathcal{R}^{C \times B \times L}$ through the hidden state of the final GRU layer. C represents the hidden unit size of GRU, B indicates the number of directions of GRU, and L denotes the number of GRU layers. We build two consecutive GRU layers of 4096 hidden units with bidirectional configuration, so $C = 4096$, $B = 2$, and $L = 2$. Thus, we can derive that:

$$E_{\text{style}} = \text{GRU}(\Delta S), \quad (3)$$

where **GRU** represents the operation of bidirectional two GRU layers embedded in StyleGRU module.

3.1.3 3D ResNet-50 module

The content feature C_{content} is derived from the feature extraction performed by the 3D CNN as:

$$C_{\text{content}} = \text{3D CNN}(I), \quad (4)$$

where **3D CNN** is an operation function for extracting a feature from 3D CNN model. Thus, C_{content} contains information about visual and temporal inconsistency. We utilize

the feature of $\mathcal{R}^{1024 \times 16}$ by applying a global average pooling to the last convolution layer of 3D-R50 architecture [3] pre-trained in FTCN [57].

3.1.4 Style attention module

We introduce a Style Attention Module (SAM) that leverages the style embedding vector obtained from the StyleGRU. To apply the attention mechanism, we designate C_{content} as the query, while E_{style} is allocated as the key and value. We employ the content feature as the query based on the intuition that these features hold the potential to identify anomalies in the style flow, driven by their encoded temporal inconsistency.

Thus, we design the query, key, and value of SAM as:

$$Q_C = \phi_q(C_{\text{content}}), \quad K_E = \phi_k(E_{\text{style}}), \quad V_E = \phi_v(E_{\text{style}}), \quad (5)$$

where ϕ_q , ϕ_k , and ϕ_v are linear projection function using fully connected layer to estimate the query, key, and value, respectively. Then, we calculate a Style Attention (SA) using the query and key as:

$$\text{SA}(Q_C, K_E) = s \left(\frac{Q_C K_E}{\sqrt{d_{K_E}}} \right), \quad (6)$$

where $s(\bullet)$ represents a softmax function.

Ultimately, we compute the result of SAM by taking the dot-product of the SA with the value component. Thus,

$$\text{SAM}(E, C) = \phi(\text{SA}(Q_C, K_E) \cdot V_E), \quad (7)$$

where ϕ is the last linear projection layer of SAM.

3.1.5 Temporal transformer encoder module.

The output from the SAM is linearly projected into 16 time steps ($n=16$) and then fed into the Temporal Transformer Encoder (TTE) module. The core of TTE is composed of standard transformer encoder blocks [49], with each block including a multi-head self-attention module and multi-layer perceptrons. The results of TTE are the input to the following classification head, yielding the final deepfake prediction through a sigmoid function σ . Thus,

$$\hat{y} = \sigma(f_{cls}(\mathbf{TTE}(C_{\text{content}} + \mathbf{SAM}(E_{\text{style}}, C_{\text{content}}))), \quad (8)$$

where f_{cls} is a function of fully-connected layer to represent the classification head. We refer to FTCN [57] for the architecture of TTE.

3.2. Training Procedure

The training procedure consists of two following stages. The first stage involves training a Gated Recurrent Unit (GRU) layers to enable encoding of the temporal style latent flow suitable for deepfake detection. In the following stage, we train the 3D CNN model for content embedding, the style attention module, and the subsequent modules for the binary classification. Fig. 3 provides a comprehensive illustration of our training procedures.

3.2.1 Stage 1: style representation learning

In stage 1, we train StyleGRU to learn the representation of style flow. We employ supervised contrastive learning [23] to produce an effective representation of style flow. For the supervised contrastive learning, we organized the style flows used as inputs to the GRU layers into three categories: anchor, positive, and negative style flows.

We utilize a different sampling strategy to build the clip for anchor, positive, and negative style flows. At first, the anchor style flow (ΔS_a) is extracted from the clip sampled from the video sequence by a sliding window sampling strategy. On the other hand, the positive style flow (ΔS_p) is obtained by using a random sampling strategy. Both the anchor and the positive style flows are extracted from the preprocessed clips with the same label. In contrast, the negative style flow (ΔS_n) is obtained from the clips with labels distinct from those of the anchor and positive clips. The negative style flow is generated by randomly employing one of the two sampling strategies.

As a batch, the anchor, positive, and negative style flows are fed into the GRU layers, and their representation is trained by using a triplet loss and a classification loss. When we define the GRU’s final hidden states from ΔS_a , ΔS_p , and ΔS_n by E_a , E_p , and E_n , respectively, the triplet loss for training the GRU layers is derived as follows:

$$L_{tri} = \max(\|E_a - E_p\|_2^2 - \|E_a - E_n\|_2^2 + \alpha, 0), \quad (9)$$

where α is a margin parameter.

Additionally, we used classification loss (L_{cls}), which is designed by Binary Cross-Entropy (BCE) loss to learn appropriate features for classification. To obtain the predicted value \hat{y} for the ground truth label y , the hidden states E of the GRU are input into the auxiliary classifier. The calculated $\text{BCE}(y, \hat{y})$ loss is utilized as L_{cls} . The auxiliary classifier consists of three fully connected layers, with a residual connection applied after the first two layers. The hidden size of each fully connected layer is kept identical at 4096.

Then, the entire loss function for training StyleGRU is as follows:

$$L = L_{tri} + \lambda L_{cls}, \quad (10)$$

where λ is a user-defined hyperparameter to determine the influence degree of L_{cls} . We minimize the loss by using Adam optimizer [25].

3.2.2 Stage 2: style attention-based deepfake detection.

In stage 2, to perform the real/fake binary classification task, we employ a conventional Binary Cross-Entropy (BCE) loss to train the 3D ResNet-50, SAM, and TTE modules, while fixing the StyleGRU module trained at stage 1. The BCE loss (L_{bce}) is applied to reduce the gap between the fake prediction \hat{y} of Eq. 8 and its corresponding ground truth label, where 0 is real and 1 is fake. We used the momentum-driven SGD optimizer for the training process of stage 2.

4. Experiments

We conduct experiments in this section to demonstrate the superiority of our algorithm. Experiments under cross-dataset and cross-manipulation conditions serve as benchmark settings to assess the generalization performance of deepfake detection models.

Implementation Details. During the preprocessing phase, we utilized Dlib [24] to extract facial landmarks and perform alignment, as well as Retinaface [10] to carry out face cropping. Following this preprocessing, the real and fake clips fed into the model consist of a total of $l = 32$ frames each. In stage 1, a dropout layer with a dropout rate of $p = 0.2$ is included between the input style flow and the StyleGRU. Furthermore, in StyleGRU, an RNN dropout rate of $p = 0.1$ is applied. As training parameters, a total of 50,000 samples are organized with a batch size of 256 over 100 epochs for contrastive learning. For optimization, the ADAM optimizer is used with a learning rate of $5e^{-4}$.

In stage 2, we use the momentum SGD optimizer and set the weight decay to $1e^{-4}$. The training is conducted with a batch size of 16. For the first 10 epochs, the learning rate of SGD gradually warms up from 0.01 to 0.1. Subsequently, a cosine learning rate scheduler is used to decrease the learning rate throughout 100 epochs.

| Method | CDF | DFD | FSh | DFo | Avg |
|---------------------|-------------|-------------|-------------|-------------|-------------|
| Xception [8] | 73.7 | - | 72.0 | 84.5 | - |
| CNN-aug [50] | 75.6 | - | 65.7 | 74.7 | - |
| PatchForensics [4] | 69.6 | - | 57.8 | 81.8 | - |
| Multi-task [35] | 75.7 | - | 66.0 | 77.7 | - |
| FWA [28] | 69.5 | - | 65.5 | 50.2 | - |
| Two-branch [34] | 76.7 | - | - | - | - |
| DCL [44] | 82.3 | 91.6 | - | - | - |
| Face X-ray [27] | 79.5 | <u>95.4</u> | 92.8 | 86.8 | 88.6 |
| NoiseDF [51] | 75.9 | - | - | 70.9 | - |
| SLADD [5] | 79.7 | - | - | - | - |
| SBI-R50 [43] | 85.7 | 94.0 | 78.2 | 91.4 | 87.3 |
| CNN-GRU [41] | 69.8 | - | 80.8 | 74.1 | - |
| STIL [15] | 75.6 | - | - | - | - |
| LipForensics [16] | 82.4 | - | 97.1 | 97.6 | - |
| RealForensics* [17] | 85.6 | 82.2 | 99.6 | 99.8 | 91.8 |
| FTCN* [57] | <u>86.9</u> | 94.4 | 98.8 | 98.8 | 94.7 |
| AltFreeze* [52] | 89.0 | 93.7 | <u>99.2</u> | <u>99.0</u> | <u>95.2</u> |
| StyleGRU | 62.5 | 72.2 | 90.7 | 77.7 | 79.8 |
| Ours | 89.0 | 96.1 | 99.0 | 99.0 | 95.8 |

Table 1. **Generalization to cross-dataset.** We present the AUC scores (%) on the CDF, DFD, FSh, and DFo datasets to assess the cross-dataset performance of our framework. The highest-performing score is marked in **bold**. The underlined values correspond to the Top-2 best methods. The asterisk (*) denotes that we have reproduced the results using officially provided weights. The results of other methods were obtained from AltFreeze [52].

Datasets and Measurement. To ensure robustness under various conditions, we used the following deepfake detection datasets: (1) **FaceForensics++** (FF++) [40] comprises 1,000 original videos and 4,000 fake videos produced in four manipulation methods, Deepfake (DF) [9], Face2Face (F2F) [46], FaceSwap (FS) [32], NeuralTexture (NT) [47], and we utilize the c23 version. (2) **FaceShifter** (FSh) [26] and (3) **DeeperForensics** [19] result from the employment of optimized manipulation techniques to the original videos of FF++. We executed our experiments using test videos adhering to the same split as in FF++. (4) **CelebDF-v2** (CDF) [29] is composed of 590 real videos and 5,639 fake videos. For our study, we engaged a subset of 518 test videos. (5) **DeepfakeDetection** (DFD) [13] consists of 363 real videos and 3071 synthetic videos.

To train our model, we used the training videos from the FF++ [40] dataset with light compression, specifically identified as c23. The evaluation metric is the Area Under the receiver operating characteristic Curve (AUC). Most deepfake detection algorithms have utilized AUC as their evaluation metric, and Lipforensics [16], which is a fundamental benchmark for video-based deepfake detection, also employed AUC.

Comparisons. We have compared both image-based and video-based deepfake detection models. Xception [40] is a method of training the Xception model [8], which is a

widely used CNN architecture. CNN-aug [50] discovered that CNN-generated images can be readily identified by a CNN model. PatchForensics [4] indicated that the use of patch-based classifiers can contribute to robust deepfake detection. Multi-task [35] approach employed an architecture similar to that of an auto-encoder for deepfake detection. FWA [28] introduced a training data generation technique that relies on deteriorating the quality of the GAN-generated source images. Two-branch [34] method utilized multi-task learning within the context of a deepfake detection dataset. SLADD [5] introduced an algorithm that automatically learns effective augmentation techniques for deepfake detection. SBI [43] proposed a novel approach to blending artifacts using self-images. DCL [44] applied contrastive learning to the task of deepfake detection in the image domain. NoiseDF [51] compared the features of the cropped face and background squares using noise to detect deepfakes. Face X-ray [27] identified face forgery by segmenting and delineating the blending border between source and target images. CNN-GRU [41] highlighted the potential of utilizing CNN and GRU to investigate temporal features. STIL [15] specifically focused on learning spatio-temporal inconsistency. LipForensics [16] exhibited robust generalization performance through the use of a trained model to capture the representation of natural mouth movement. RealForensics [17] improved the ability to generalize using multimodal architecture. FTCN [57] argued that leveraging temporal inconsistency is beneficial for enhancing the generalizability of deepfake detection algorithms. AltFreeze [52] is an architecture that demonstrates good generalization performance by separately training spatial information and temporal information.

4.1. Generalization to cross-datasets

In the context of our framework, we have set the cross-dataset performance as a proxy for the generalizability in real-world situations. We train the proposed framework on the FF++ and then evaluate its performance on the CDF, DFD, FSh, and DFo datasets. Table 1 shows that our framework exhibits the best detection performance in the CDF, DFD, and DFo datasets, which include high-quality deepfake videos. The performance on the FSh dataset also reaches a top-2 level, with the highest average score, indicating that our model is a deepfake detection algorithm with generalization capability. ‘StyleGRU’ represents the cross-dataset score using the classification layer directly connected after the StyleGRU for the Stage 1 training. The score of StyleGRU is superior to that of CNN-GRU [41], which is a similar structure that utilizes image features, on the FSh and DFo datasets. This suggests the potential for a better impact on generalization performance by utilizing the flow of style latent vectors rather than simply relying on the image-based feature.

| Method | Train on remaining three | | | | Avg |
|--------------------|--------------------------|-------------|-------------|-------------|-------------|
| | DF | FS | F2F | NT | |
| Xception [8] | 93.9 | 51.2 | 86.8 | 79.7 | 77.9 |
| CNN-aug [50] | 87.5 | 56.3 | 80.1 | 67.8 | 72.9 |
| PatchForensics [4] | 94.0 | 60.5 | 87.3 | 84.8 | 81.6 |
| Face X-ray [27] | <u>99.5</u> | 93.2 | 94.5 | 92.5 | 94.9 |
| CNN-GRU [41] | 97.6 | 47.6 | 85.8 | 86.6 | 79.4 |
| LipForensics [16] | <u>99.7</u> | 90.1 | <u>98.8</u> | <u>98.3</u> | 97.1 |
| FTCN* [57] | <u>99.8</u> | <u>99.3</u> | 95.9 | 95.3 | <u>97.5</u> |
| AltFreeze [52] | <u>99.8</u> | <u>99.7</u> | <u>98.6</u> | <u>96.2</u> | <u>98.6</u> |
| StyleGRU | 94.0 | 68.5 | 88.8 | 80.6 | 83.0 |
| Ours | <u>99.7</u> | <u>98.8</u> | <u>98.6</u> | <u>96.4</u> | <u>98.4</u> |

Table 2. **Generalization to cross-manipulation.** In order to quantitatively assess the performance of our method in cross-manipulation scenarios, we utilized video-level AUC. The asterisk (*) denotes results that we have reproduced. The underlined values correspond to the Top-3 best methods. The results of other methods were obtained from AltFreeze [52].

4.2. Generalization to cross-manipulations

For deepfake detection tasks, the evaluation of cross-manipulation performance targets fake videos generated from the same original video but using a different scheme. A common approach is to select three out of the four manipulations provided by FF++, and then the model trained by the three chosen manipulation sets is evaluated by one remaining manipulation set. DF [9] and FS [32] represent identity swap manipulation methods, while F2F [46] and NT [47] denote expression swap manipulation methods. AltFreeze [52] does not officially provide weights for this experiment setting. Therefore, it is not possible for us to reproduce. Table 2 presents a comparison of the performance of the cross-manipulation between our model and various models. For expression swap manipulation methods like F2F and NT, they exhibit relatively lower performance compared to Lipforensics [16]. Expression swap inherently lead to inconsistencies in the mouth region, which is why LipForensics appears to perform well. However, our framework displays outstanding performance across all manipulations.

4.3. Ablation Study

In this section, we conduct ablation studies to verify the approach used in our framework. we validate the necessity of the proposed StyleGRU and SAM and conduct an ablation study on the losses used in Stage 1.

4.3.1 Component-wise Ablation Study

As shown in Table 3, we conducted experiments using only the StyleGRU by feeding the hidden states of StyleGRU into an MLP classification head to detect the real/fake of input video. Even though only the classification through

| Method | Dataset | | | | | |
|--------|----------|-----|-----|----------|------|------|
| | StyleGRU | SAM | TTE | Δ | CDF | FSh |
| ✓ | - | - | ✓ | 62.5 | 90.7 | 76.6 |
| - | ✓ | ✓ | ✓ | 85.7 | 98.4 | 92.1 |
| ✓ | - | ✓ | ✓ | 87.8 | 98.3 | 93.1 |
| ✓ | ✓ | ✓ | - | 88.4 | 98.4 | 93.4 |
| ✓ | ✓ | ✓ | ✓ | 89.0 | 99.0 | 94.0 |

Table 3. **Framework ablation.** We evaluate our framework on the CDF and FSh datasets using video-level AUC scores and conduct ablation studies for each stage of the proposed methodology. All models utilized the FF++ dataset for training.

| Loss | | Dataset | | |
|-----------|-----------|---------|------|------|
| L_{cls} | L_{tri} | FF++ | CDF | FSh |
| ✓ | - | 99.2 | 88.6 | 99.0 |
| - | ✓ | 98.8 | 88.4 | 98.8 |
| ✓ | ✓ | 99.1 | 89.0 | 99.0 |

Table 4. **Stage 1 loss ablation.** We conduct an ablation study for the losses employed in stage 1. Each experiment uses the different combinations of training losses for StyleGRU, while preserving the training process of stage 2. The model is trained by the FF++ dataset, and evaluated by video-level AUC.

StyleGRU was conducted, ours still demonstrates high scores on the Fsh dataset. In addition, considering the significant performance drop when StyleGRU is not used, it can be observed that our proposed StyleGRU contributes to improving cross-dataset performance. The influence of the proposed SAM in Stage 2 on the outcome is also evaluated. For experiments without StyleGRU, we processed style latent vectors extracted from clips using an MLP to match the dimension of the 3D CNN features. In the case of experiments without the SAM, we simply added the StyleGRU feature to the 3D CNN feature. The absence of the SAM during the blending of content and style features leads to a substantial decline in performance. This quantitatively demonstrates the efficacy of the SAM trained during stage 2 in achieving the adapted feature for robust deepfake detection. Finally, the influence of differentiating (Δ) the style latent vectors is investigated by replacing the input of the style flow with the original style latent vectors. During the overall stage 1 training process, it was observed that the lack of differencing in style latent vectors as a pre-processing to be fed into the StyleGRU leads to a slower decrease in the loss. Training without differentiation was confirmed to not only lead to slower progression in stage 1 learning but also adversely affect cross-dataset performance.

4.3.2 Ablation Study for Losses

Ablation studies are conducted on various options of training losses that can be employed at stage 1 training for StyleGRU representation learning. As shown in Table 4, when

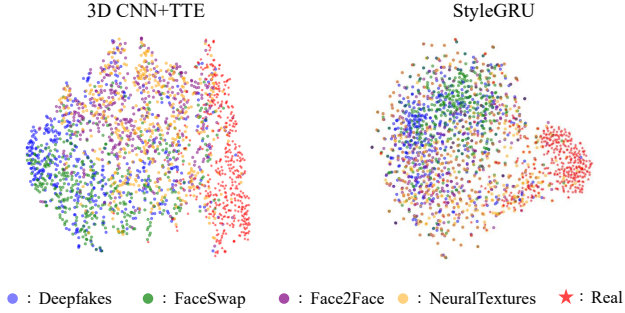


Figure 4. **t-SNE visualization.** The t-SNE visualization was conducted using the final layer features of a model trained on the FF++ test set. For hyperparameters, we utilized 1000 iterations with a perplexity of 40 and PCA to 30.

only classification loss (L_{cls}) is used, our model exhibits best performance on the training dataset, while suffering from the reduction of domain generality. Using only triplet loss (L_{tri}) results in relatively reduced performance due to the lack of distinctiveness between the real and fake videos for classification. Although employing both L_{cls} and L_{tri} results in a slight performance decrease on the seen dataset, we can see that their concurrent employment proves to be beneficial on the unseen datasets.

4.4. Analytic Experiments

In this section, we perform visualizations for further analysis of the proposed style latent flow in our study. We also analyze the perturbation robustness of our algorithm, which is important to consider in real-world scenarios.

4.4.1 t-SNE for Domain Generalization

To evaluate the learning performance of StyleGRU, we conduct t-SNE visualizations on the FF++ test set using a model trained on the FF++ dataset. As illustrated in Fig. 4, despite StyleGRU utilizing significantly fewer parameters, our model manages to extract more distinct features in the feature space. This characteristic represents our intuition that the distinctiveness of style latent vectors aligns well with the generality of deepfake video detection tasks across various generation algorithms.

4.4.2 Robustness to Unseen Perturbation

To evaluate the robustness of perturbation, experiments under the conditions proposed by DFo [19] are conducted to observe performance changes. In Fig. 5, performance under specific perturbation conditions exhibits a slight deficiency, which is attributed to the fact that the pSp encoder [38] used for extracting style latent vectors was created without considering noise. However, our model demonstrates a high level of robustness against most perturbations.




Figure 5. **Robustness to Unseen Perturbation.** The performance changes based on the video-level AUC metric when applying two perturbations at five different degradation levels. The perturbation follows the approach provided by DeeperForensics [19].

5. Conclusion

In this paper, we first demonstrate that the style latent flow extracted from consecutive frames of a video, serves as an effective cue for deepfake detection. Additionally, we confirm that leveraging an attention mechanism based on style flow facilitates model generalization. The style flow feature is trained via a supervised-contrastive representation learning process, and we develop a classification model using the Style Attention Module. The superior performance of the proposed algorithm validates the importance of temporal changes in style latent vectors to generalize the deepfake video detection task. While this paper focused on the facial attributes extracted by StyleGAN, our future plans is to broaden the scope of attributes for deepfake video detection to encompass various objects. We anticipate that utilizing StyleGAN pre-trained on other subjects, such as images of dogs and scenery, could broaden the applicability of our method to a wider range of deepfake subjects.

The model we propose includes an additional step of extracting style latent vectors through the pSp encoder. As a result, the data preprocessing takes a relatively long time. However, we confirms the positive impact of using style flow on the generalization performance of deepfake video detection models and provides a new direction for addressing synthesized videos created through GAN inversion.

Acknowledgements: This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (2021-0-01341, Artificial Intelligence Graduate School Program(Chung-Ang University), 2020-0-01336, AIGS program (UNIST), and 2021-0-01778 Development of human image synthesis and discrimination technology below the perceptual threshold) and Culture, Sports and Tourism R&D Program through the Korea Creative Content Agency grant funded by the Ministry of Culture, Sports and Tourism in 2024 (Project Name : Development of high-freedom large-scale user interaction technology using multiple projection spaces to overcome low-light lighting environments. , Project Number : RS-2023-00222280).

References

- [1] Yuval Alaluf, Or Patashnik, and Daniel Cohen-Or. Restyle: A residual-based stylegan encoder via iterative refinement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6711–6720, 2021. 2
- [2] Weiming Bai, Yufan Liu, Zhipeng Zhang, Bing Li, and Weiming Hu. Aunet: Learning relations between action units for face forgery detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24709–24719, 2023. 2
- [3] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 4
- [4] Lucy Chai, David Bau, Ser-Nam Lim, and Phillip Isola. What makes fake images detectable? understanding properties that generalize. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI 16*, pages 103–120. Springer, 2020. 6, 7
- [5] 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*, pages 18710–18719, 2022. 2, 6
- [6] Shen Chen, Taiping Yao, Yang Chen, Shouhong Ding, Jilin Li, and Rongrong Ji. Local relation learning for face forgery detection. In *Proceedings of the AAAI conference on artificial intelligence*, pages 1081–1088, 2021. 2
- [7] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014. 4
- [8] 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. 6, 7
- [9] Deepfakes. faceswap. <https://github.com/deepfakes/faceswap>, 2019. [Accessed 27-07-2023]. 6, 7
- [10] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotzia, and Stefanos Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5203–5212, 2020. 5
- [11] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 1
- [12] 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*, pages 3994–4004, 2023. 2
- [13] Nick Dufour and Andrew Gully. Contributing Data to Deepfake Detection Research — ai.googleblog.com. <https://ai.googleblog.com/2019/09/contributing-data-to-deepfake-detection.html>, 2019. [Accessed 30-07-2023]. 6
- [14] Qiqi Gu, Shen Chen, Taiping Yao, Yang Chen, Shouhong Ding, and Ran Yi. Exploiting fine-grained face forgery clues via progressive enhancement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 735–743, 2022. 1, 2
- [15] Zhihao Gu, Yang Chen, Taiping Yao, Shouhong Ding, Jilin Li, Feiyue Huang, and Lizhuang Ma. Spatiotemporal inconsistency learning for deepfake video detection. In *Proceedings of the 29th ACM international conference on multimedia*, pages 3473–3481, 2021. 1, 6
- [16] Alexandros Haliassos, Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Lips don’t lie: A generalisable and robust approach to face forgery detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5039–5049, 2021. 1, 2, 6, 7
- [17] 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*, pages 14950–14962, 2022. 6
- [18] Yonghyun Jeong, Doyeon Kim, Youngmin Ro, and Jongwon Choi. Freggan: robust deepfake detection using frequency-level perturbations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1060–1068, 2022. 2
- [19] Liming Jiang, Ren Li, Wayne Wu, Chen Qian, and Chen Change Loy. Deepforensics-1.0: A large-scale dataset for real-world face forgery detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2889–2898, 2020. 1, 6, 8
- [20] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 2, 3
- [21] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. 2, 3
- [22] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34:852–863, 2021. 2
- [23] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673, 2020. 2, 5
- [24] Davis E King. Dlib-ml: A machine learning toolkit. *The Journal of Machine Learning Research*, 10:1755–1758, 2009. 5
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [26] 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. 6

- [27] 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. 2, 6, 7
- [28] Yuezun Li and Siwei Lyu. Exposing deepfake videos by detecting face warping artifacts. arxiv 2018. *arXiv preprint arXiv:1811.00656*, 2018. 1, 6
- [29] 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. 1, 6
- [30] Zhengzhe Liu, Xiaojuan Qi, and Philip H.S. Torr. Global texture enhancement for fake face detection in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [31] Zhian Liu, Maomao Li, Yong Zhang, Cairong Wang, Qi Zhang, Jue Wang, and Yongwei Nie. Fine-grained face swapping via regional gan inversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8578–8587, 2023. 3
- [32] MarekKowalski. Faceswap. <https://github.com/MarekKowalski/FaceSwap>, 2019. [Accessed 30-07-2023]. 6, 7
- [33] Francesco Marra, Diego Gagnaniello, Luisa Verdoliva, and Giovanni Poggi. Do gans leave artificial fingerprints? In *2019 IEEE conference on multimedia information processing and retrieval (MIPR)*, pages 506–511. IEEE, 2019. 1
- [34] Iacopo Masi, Aditya Killekar, Royston Marian Mascarenhas, Shenoy Pratik Gurudatt, and Wael AbdAlmageed. Two-branch recurrent network for isolating deepfakes in videos. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VII 16*, pages 667–684. Springer, 2020. 6
- [35] Huy H Nguyen, Fuming Fang, Junichi Yamagishi, and Isao Echizen. Multi-task learning for detecting and segmenting manipulated facial images and videos. In *2019 IEEE 10th international conference on biometrics theory, applications and systems (BTAS)*, pages 1–8. IEEE, 2019. 6
- [36] Utkarsh Ojha, Yuheng Li, and Yong Jae Lee. Towards universal fake image detectors that generalize across generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24480–24489, 2023. 1
- [37] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021. 1
- [38] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2287–2296, 2021. 2, 3, 8
- [39] Jonas Ricker, Simon Damm, Thorsten Holz, and Asja Fischer. Towards the detection of diffusion model deepfakes. *arXiv preprint arXiv:2210.14571*, 2022. 1
- [40] 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. 1, 2, 6
- [41] Ekraam Sabir, Jiaxin Cheng, Ayush Jaiswal, Wael AbdAlmageed, Iacopo Masi, and Prem Natarajan. Recurrent convolutional strategies for face manipulation detection in videos. *Interfaces (GUI)*, 3(1):80–87, 2019. 1, 2, 6, 7
- [42] Ekraam Sabir, Jiaxin Cheng, Ayush Jaiswal, Wael AbdAlmageed, Iacopo Masi, and Prem Natarajan. Recurrent convolutional strategies for face manipulation detection in videos. *Interfaces (GUI)*, 3(1):80–87, 2019. 2
- [43] Kaede Shiohara and Toshihiko Yamasaki. Detecting deepfakes with self-blended images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18720–18729, 2022. 2, 6
- [44] Ke Sun, Taiping Yao, Shen Chen, Shouhong Ding, Jilin Li, and Rongrong Ji. Dual contrastive learning for general face forgery detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2316–2324, 2022. 6
- [45] Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhofer, and Christian Theobalt. Stylerig: Rigging stylegan for 3d control over portrait images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6142–6151, 2020. 3
- [46] 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. 6, 7
- [47] Justus Thies, Michael Zollhofer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *Acm Transactions on Graphics (TOG)*, 38(4):1–12, 2019. 6, 7
- [48] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics (TOG)*, 40(4):1–14, 2021. 2
- [49] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 5
- [50] 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*, pages 8695–8704, 2020. 6, 7
- [51] Tianyi Wang and Kam Pui Chow. Noise based deepfake detection via multi-head relative-interaction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 14548–14556, 2023. 1, 6
- [52] Zhendong Wang, Jianmin Bao, Wengang Zhou, Weilun Wang, and Houqiang Li. Altfreezing for more general video

- face forgery detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4129–4138, 2023. [2](#), [6](#), [7](#)
- [53] Chao Xu, Jiangning Zhang, Miao Hua, Qian He, Zili Yi, and Yong Liu. Region-aware face swapping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7632–7641, 2022. [3](#)
- [54] Yangyang Xu, Bailin Deng, Junle Wang, Yanqing Jing, Jia Pan, and Shengfeng He. High-resolution face swapping via latent semantics disentanglement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7642–7651, 2022. [3](#)
- [55] Zhiyuan Yan, Yong Zhang, Yanbo Fan, and Baoyuan Wu. UCF: uncovering common features for generalizable deepfake detection. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, pages 22355–22366. IEEE, 2023. [2](#)
- [56] Hanqing Zhao, Tianyi Wei, Wenbo Zhou, Weiming Zhang, Dongdong Chen, and Nenghai Yu. Multi-attentional deepfake detection. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2185–2194, 2021. [2](#)
- [57] Yinglin Zheng, Jianmin Bao, Dong Chen, Ming Zeng, and Fang Wen. Exploring temporal coherence for more general video face forgery detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 15044–15054, 2021. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#)
- [58] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *European conference on computer vision*, pages 592–608. Springer, 2020. [2](#)
- [59] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A Efros. Generative visual manipulation on the natural image manifold. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part V 14*, pages 597–613. Springer, 2016. [2](#)
- [60] Yuhao Zhu, Qi Li, Jian Wang, Cheng-Zhong Xu, and Zhenan Sun. One shot face swapping on megapixels. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4834–4844, 2021. [3](#)