

Undercover Deepfakes: Detecting Fake Segments in Videos

Sanjay Saha¹, Rashindrie Perera^{*2}, Sachith Seneviratne^{*3}, Tamasha Malepathirana², Sanka Rasnayaka¹, Deshani Geethika², Terence Sim¹, and Saman Halgamuge³

¹National University of Singapore, Singapore, ^{2,3}University of Melbourne, Australia,

¹{sanjay-s,sanka,tsim}@comp.nus.edu.sg,

²{cdperera,tmalepathira,dpoddenige}@student.unimelb.edu.au,


³{sachith.seneviratne,saman}@unimelb.edu.au

Abstract

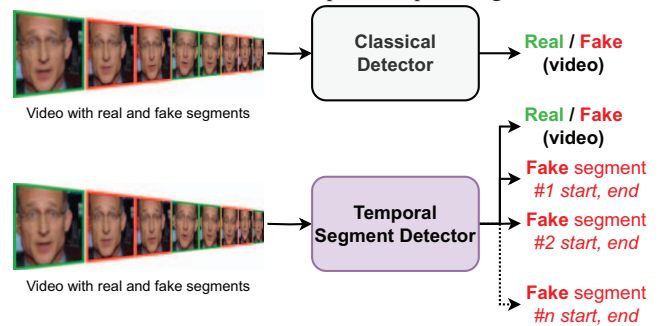
The recent renaissance in generative models, driven primarily by the advent of diffusion models and iterative improvement in GAN methods, has enabled many creative applications. However, each advancement is also accompanied by a rise in the potential for misuse. In the arena of the deepfake generation, this is a key societal issue. In particular, the ability to modify segments of videos using such generative techniques creates a new paradigm of deepfakes which are mostly real videos altered slightly to distort the truth. This paradigm has been under-explored by the current deepfake detection methods in the academic literature. In this paper, we present a deepfake detection method that can address this issue by performing deepfake prediction at the frame and video levels. To facilitate testing our method, we prepared a new benchmark dataset where videos have both real and fake frame sequences with very subtle transitions. We provide a benchmark on the proposed dataset with our detection method which utilizes the Vision Transformer based on Scaling and Shifting [38] to learn spatial features, and a Timeseries Transformer to learn temporal features of the videos to help facilitate the interpretation of possible deepfakes. Extensive experiments on a variety of deepfake generation methods show excellent results by the proposed method on temporal segmentation and classical video-level predictions as well. In particular, the paradigm we address will form a powerful tool for the moderation of deepfakes, where human oversight can be better targeted to the parts of videos suspected of being deepfakes. All experiments can be reproduced at: github.com/rgb91/temporal-deepfake-segmentation.

*These authors contributed equally to this work.

1. Introduction



(a) Simplified illustration of sample videos from the newly introduced benchmark dataset for temporal deepfake segment detection.



(b) Comparison of our method with the classical deepfake detection. Figure 1: (a) We propose a new deepfake benchmark dataset consisting of videos with one or two manipulated segments, represented in the two images respectively. Fake frames are indicated by smaller red boxes and genuine video frames are denoted by green borders. (b) Our proposed deepfake detection method employs temporal segmentation to classify video frames as real or fake and identify the intervals containing the manipulated content. This is a departure from the conventional binary classification of videos as either entirely genuine or entirely manipulated.

Deep learning has made significant advances over the last few years, with varying degrees of societal impact. The advent of diffusion models as viable alternatives to hitherto established generative models has revolutionized the domains of textual/language learning, visual generation, and

cross-modal transformations [54]. This, alongside other recent advancements in generative AI such as GPT4 [6] has also brought to attention the social repercussions of advanced AI systems capable of realistic content generation.

There exist many methods to tackle the deepfake detection problem formulated as a binary classification problem [1, 5, 13]. A common pitfall of these methods is the inability to generalize to unseen deepfake creation methods. However, there is a more pressing drawback in these studies that we aim to highlight and address in this paper. Considering the social impact of a deepfake video, we hypothesize that rather than fabricating an entire fake video, a person with malicious intent would alter smaller portions of a video to misrepresent a person’s views, ideology and public image. For example, an attacker can generate a few fake frames to replace some real frames in a political speech, thus distorting their political views which can lead to considerable controversy and infamy.

The task of identifying deepfake alterations within a longer video, known as deepfake video temporal segmentation, is currently not well explored or understood. These types of deepfakes pose a more difficult challenge for automated deepfake analysis compared to other types of deepfakes. Moreover, they also pose a greater threat to society since the majority of the video may be legitimate, making it appear more realistic and convincing. Additionally, these deepfakes require significant manual oversight, especially in the moderation of online content platforms, as identifying the legitimacy of the entire video requires manual interpretation. However, performing frame-level detection allows the human moderator to save time by focusing only on the fake segments.

Figure 1(a) demonstrates deepfake videos where the entire video is not fake, but some of the real frames were replaced by fake frames. We present a benchmark dataset with videos similar to those in Figure 1(a) to test our method on the temporal deepfake segmentation problem. In the temporal deepfake segmentation problem the detector makes frame level predictions and calculates the start and end of the fake sequences i.e. fake-segments. This differs from classical deepfake detection where the detector makes a video level prediction as demonstrated in Figure 1(b).

Problem Definition

Deepfake Temporal Segmentation task is defined as,

Given an input video identify temporal segments within the video that are computer generated i.e. fakes. The output of this task is a labeling of ‘real’ or ‘fake’ for each frame, which we call a temporal segmentation map. We can frame the classical deepfake detection problem as a special case of the temporal segmentation task, in which ALL frames are labeled either as ‘real’ or ‘fake’. With an emphasis on the novel deepfake temporal segmentation task, this paper makes the following contributions,

- We emphasize on the new threat of faking small parts of a longer video to pass it off as real. Current detection methods ignore this threat, since they assume the entire video is real or fake. This can be addressed through proposed temporal segmentation of videos. This provides a new direction for future research.
- We curated a new dataset specifically for deepfake temporal segmentation, which will be publicly available for researchers to evaluate their methods. Our rigorous experiments establish benchmark results for temporal segmentation of deepfakes, providing a baseline for future work.

2. Related Work

Face Image Synthesis: Manipulation of face images has always been a popular research topic in the media forensics [66, 52] and biometrics domain. Synthesized digital faces can be used to deceive humans as well as machines and software. Prior to deepfakes, digitally manipulated faces [56, 57, 30] were utilized mainly to fool biometric verification and identification methods e.g., face recognition systems. Consequently, deepfake methods [27, 16, 15, 72] started to generate very realistic fake videos of faces and became much more popular as a result. This led to a series of research works on developing a number of deepfake generation methods, categorized into mainly two types: Face swapping [17, 33, 47] and Face reenactment [63, 61].

Deepfake generation methods have since improved significantly by advancing existing methods and better software integration, as in Deepfacelab [47]. This has helped creators of deepfakes to create longer videos, including seamlessly blending fake frames with real frames, which allows one to have both real and fake video segments within the same deepfake video. Through more recent developments in generative AI [54, 64, 45] we are at the brink of experiencing even higher quality and more subtle deepfakes, raising the need for updated research in this area.

Deepfake Detection: Initial works on deepfake detection methods [59, 51, 69, 46, 2] focused on detecting artifacts in deepfaked face images, such as irregular eye colors, asymmetric blinking eyes, abnormal heart beats, irregular lip, mouth and head movements [36, 60, 12, 42]. Some other earlier works tried to find higher-level variability in the videos: erroneous blending after face swaps, or identity-aware detection approach [68, 21, 34, 14]. Compared to these earlier works, more recent studies [44, 50, 3, 73, 4] that are independent of artifact-based detection have achieved astounding results in detecting fake videos from most of the state-of-the-art datasets. Recently, more works [8, 67, 70, 28, 43, 31, 26, 23, 71, 58] have given increased attention towards generalizability of the detectors to detect deepfakes from unseen methods.

Temporal Segmentation: Although deepfake detection methods have seen significant progress in recent years, only a few studies [25, 11, 7, 24] have looked into the problem of temporal segmentation task where in a long video only one or more short segments is altered while the rest of the frames are real. In this paper, we introduce a new, easily reproducible dataset based on the FaceForensics++ [55] and a method for not only detecting deepfake videos but also segmenting the fake frame-segments within them. The proposed approach can accurately identify one or more fake segments in a deepfake video, which can mitigate the risks associated with deepfakes that remain well-blended within real frames in a video.

3. Methodology

We propose a two-stage method as shown in Figure 2. The first stage employs a Vision Transformer (ViT) based on Scaling and Shifting (SSF) [38] to extract the frame-level features of the videos. Specifically, ViT learns a single vector representation for each frame and these feature vectors are sequentially accumulated and windowed using the sliding window technique. The TsT is an adaptation of the original transformer encoder from [65]. TsT learns the temporal features from the learned ViT features and uses them for classification. The feature vectors from stage 1 are sequentially accumulated and windowed using the sliding window technique. It is important to use sequentially windowed feature vectors as input to the TsT since we need the temporal features to be learned for better temporal segmentation.

3.1. Model architecture

3.1.1 Vision Transformer (ViT) and Scaling and Shifting (SSF)

ViTs have achieved state-of-the-art results on several image classification benchmarks, demonstrating their effectiveness as an alternative to convolutional neural networks (CNNs). The employed ViT model first partitions the input image $I \in \mathbb{R}^{H \times W \times C}$ into a set of smaller patches of size $N \times N$ where H , W , C and N correspond to the height, width, number of channels of the image, and the height and width of each patch, respectively. Each patch is then represented by a d -dimensional feature vector, which is obtained by flattening the patch into a vector of size $N^2 C$ and applying a linear projection to reduce its dimensionality. Next, to allow the model to learn the spatial relationships between the patches, positional encodings are added to the patch embeddings. The resulting patch embeddings are concatenated together to form a sequence, and a learnable class embedding that represents the classification output is prepended to the sequence which is then input through a series of transformer layers. Each transformer layer consists of a multi-head self-attention mechanism, which allows the model to

attend to different parts of the input patches, a multi-layer perceptron (MLP), and a layer normalization (Fig. 4). Finally, a classification head is attached at the end of the transformer layers, which produces a probability distribution over the target classes.

Recently, there has been an upsurge in the use of parameter-efficient fine-tuning methods [38, 9] to fine-tune only a smaller subset of parameters in large pre-trained models such as ViTs, leading to better performance in downstream tasks compared to conventional end-to-end fine-tuning and linear probing. We use one such method, called SSF [38] to fine-tune the pre-trained ViT model used in our pipeline (Fig. 2). SSF attempts to alleviate the distribution mismatch between the pre-trained task and the downstream deep fake feature extraction task by modulating deep features. Specifically, during the fine-tuning phase, the original network parameters are frozen, and SSF parameters are introduced at each operation to learn a linear transformation of the features, as shown in Fig. 4.

As done in the original work, we too insert SSF parameters after each operation including multi-head self-attention, MLP, layer normalization, etc. Specifically, given the input $x \in \mathbb{R}^{N^2+1} \times d$, the output $y \in \mathbb{R}^{N^2+1} \times d$ (also the input to the next operation) is calculated by

$$y = \gamma \cdot x + \beta \quad (1)$$

where $\gamma \in \mathbb{R}^d$ and $\beta \in \mathbb{R}^d$ are the scale and shift parameters, respectively.

3.1.2 Timeseries Transformer (TsT)

Architecture and training. The employed TsT is an adaptation of the sequence to sequence transformer in [65]. The transformer architecture is designed to learn and classify from sequential data instead of generating another sequence. It is composed of multiple transformer blocks and an MLP head. Each transformer block has a multi-head attention mechanism and a feed-forward block as shown in Figure 4(c).

Our method generates frame-level predictions for the input videos, which may contain some noisy predictions. To address this issue, we used a simple smoothing technique based on majority voting over a sliding window of size 15. It takes a majority vote from the predictions of the frames within the window around a particular frame. By smoothing out the noisy predictions, our approach improves performance, as demonstrated in Table 6.

Data processing. For the TsT, we accumulated the feature vectors from the ViT sequentially for each video and split them into overlapping windows of size W . That is, we have features of W sequential frames in one window.

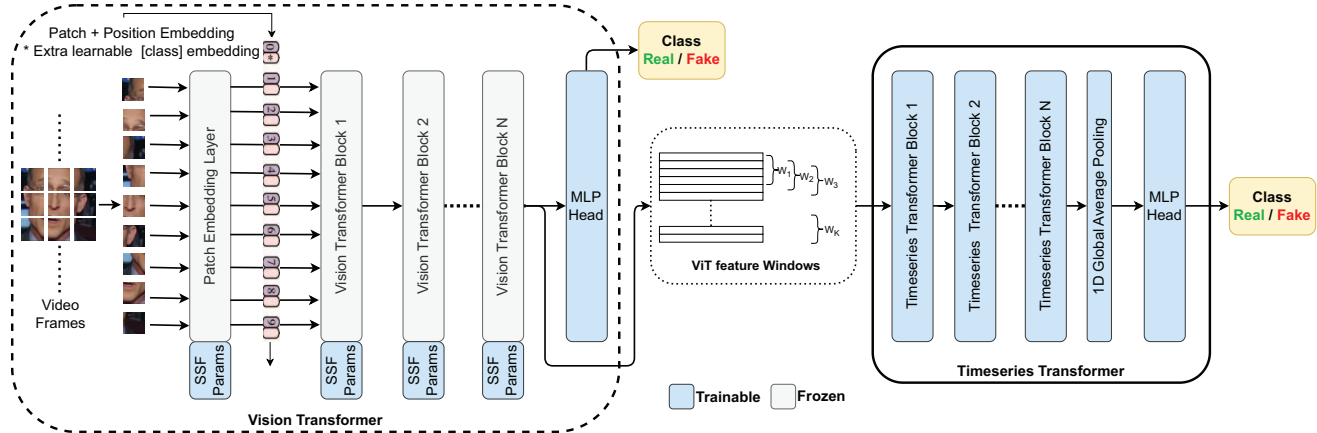


Figure 2: Our proposed detection method’s model architecture comprises two main blocks: the Vision Transformer (ViT) and the Timeseries Transformer (TsT). The ViT fine-tunes a pre-trained model for deepfake detection using the Scaling and Shifting (SSF) method, learning spatial features. On the other hand, the TsT focuses on temporal features. The ViT’s spatial features are sequentially accumulated and split into overlapping windows before inputting into the TsT.

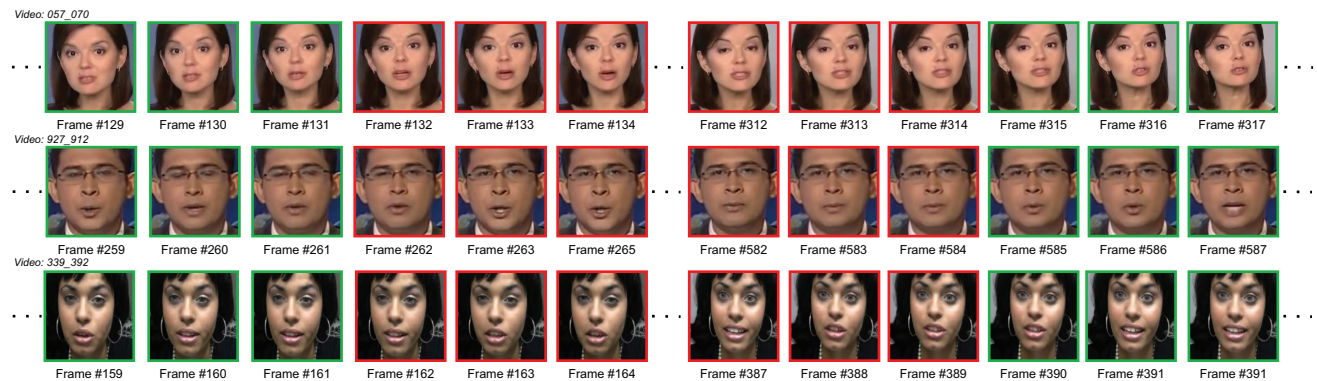


Figure 3: Samples from temporal dataset where segments were carefully selected (hand-crafted) through manual inspection. In this illustration, we focus specifically on where the transitions takes place. This alteration from **real** to **fake** frames and vice versa are subtle. The videos start with a real sequence and subtly changes to a fake sequence before going back to real sequence, and they can easily deceive human inspection.

3.2. Temporal Deepfake Segment Benchmark

Nearly all deepfake related studies have assumed that all the frames in deepfake videos are from one class, i.e. either fake or real. However, it is possible to have both types of frames within one video. With sophisticated deepfake generation methods such as Neural Textures [62], Deepfacelab [47] etc., it is possible to create fake segments that blend masterfully with the neighboring real segments within a video. This makes it very hard even for experienced human eyes to detect and to separate the fake segments from the real ones. Hence, it is important to explore automated temporal segmentation of deepfake videos.

3.2.1 Dataset

The existing deepfake datasets contain videos where all the frames in a video are from one class: ‘real’ or ‘fake’. To the best of our knowledge, there are currently no publicly

available datasets crafted specially for the temporal segmentation task. Therefore, we created the first benchmark dataset with videos that contain fake and real frames to test temporal segmentation. The dataset was created from a re-sampled subset from the original FaceForensics++ (FF++) dataset [55]. There are five face manipulation techniques within FF++ which we refer to as sub-datasets in our paper. Each sub-dataset has 1,000 videos in total. These five sub-datasets are prepared based on five different deepfake generation methods: Deepfakes (DF) [17], FaceShifter (FSh) [33], Face2Face (F2F) [63], NeuralTextures (NT) [61], FaceSwap (FS) [32]. Since the distribution of any copy of the original dataset is limited, we will publish the code and necessary files to regenerate the temporal dataset instead of publishing the videos.

Our dataset comes in two main parts: 1) videos with hand-crafted fake segments and 2) videos with randomly

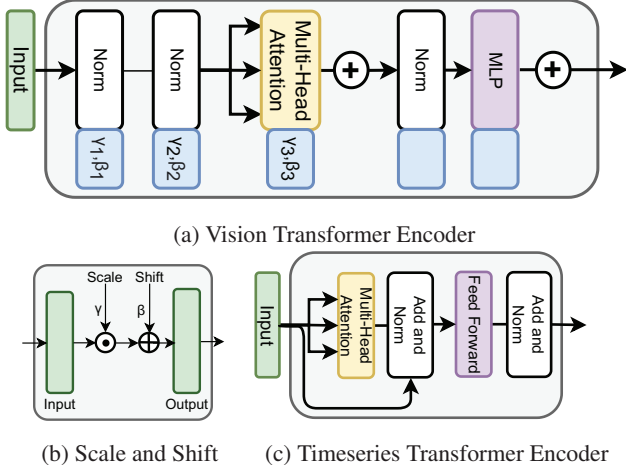


Figure 4: Architecture of the encoder blocks and Scale and Shift block.

chosen fake segments.

Hand-crafted fake-segments: Since transitioning from real to fake frames and vice versa can cause heavy temporal artifacts, it makes the videos unrealistic if the transition is not subtle. Only videos of *Neural Textures* and *Face2Face* were used to create this part of the dataset, as these methods provide the opportunity for a seamless transition. We manually selected the fake segments to make sure that the changes in sequences are most realistic. Some samples of the NT videos are visualized in Figure 3. The subtlety of the alteration even fools the careful human eyes. There are 100 videos in each sub-dataset (NT and F2F) where each video contains one fake segment.

Randomly chosen fake-segments: This part contains all the five subdatasets from FF++. Each sub-dataset contains randomly selected 100 videos. We have 500 videos in total with one fake segment, and another 500 videos with two fake segments. For videos with one fake segment, we have selected a random starting point in the first half of the video and a random choice from 125, 150, and 175 frames for the length of the fake segment. A similar strategy is also selected for videos with two fake segments. Here, the first fake segment starts at a random position within the first 125 frames and the second fake segment starts at a random position within the first 75 frames in the second half of the video. The lengths of the fake segments here are randomly chosen.

On average, 24.3% of frames were fake in the videos with one fake segment and the ratio is 41.1% for videos with two fake segments. Average length (number of frames) of a video is 633.9. Detailed break-down of these ratio and the length of the videos for each deepfake generation method in the dataset are reported in Table 7.

Model trained on	F2F		NT	
	IoU	AUC	IoU	AUC
Deepfakes (DF, ours)	0.948	0.964	0.684	0.744
Face Shifter (FSh, ours)	0.943	0.962	0.637	0.699
Face2Face (F2F, ours)	0.980	0.987	0.738	0.794
Neural Textures (NT, ours)	0.943	0.970	0.931	0.960
Face Swap (FS, ours)	0.954	0.969	0.553	0.607
FF++ (CADDM [19])	0.942	0.989	0.756	0.943
FF++ (ours)	<i>0.970</i>	0.984	<i>0.930</i>	<i>0.953</i>

Table 1: Results for temporal segmentation on the proposed temporal dataset with **hand-crafted fake-segments**. Each row indicates results from a model trained on a specific training sub-dataset; we have trained models with FaceForensics++ (FF++) and the five sub-datasets within FF++. We compare our results with CADDM[19], as shown in the second last row. The last row presents the results for the model trained on the full FF++ dataset. The columns represent the data we have tested our models on; we have tested the models on the two sub-datasets from the hand-crafted temporal segments: NT and F2F. We report IoU and AUC metrics, where the best value in a column is represented in **bold**, and the second-best value is represented in *italic*.

3.2.2 Evaluation: Intersection over Union (IoU)

Intersection over Union (IoU) is proposed to evaluate the temporal segmentation map. This metric is most commonly used to evaluate the fit of object detection bounding boxes [53, 49]. 1-D variations of IoU has been adopted for time series segment analysis, which we will be utilizing.

Let the ground truth map be $GT_{map} = \{RRRRRRFFFRR\dots\}$ and predicted segmentation map be $P_{map} = \{RRRRRRFFFRR\dots\}$. Both are 1-D vectors of equal length with a predicted Boolean class (R or F) for each frame in the video.

$$IoU = \frac{Intersection}{Union} = \frac{|GT_{map} \cap P_{map}|}{|GT_{map} \cup P_{map}|} \quad (2)$$

IoU falls in the range $[0, 1]$; where the greater the value, the better the predicted segment map. Although the theoretical lower bound of IoU is zero, in practice it is useful to understand how a random guessing algorithm will be scored. For a random guessing algorithm with probability $p = 0.5$ for each class in a binary classification problem, we have $IoU = 1/3$. This will be the random guessing baseline for IoU in our context.

4. Results

4.1. Experimental settings

Dataset. Based on the recent deepfake detection methods we have used FaceForensics++ [55] (FF++, see Section 3.2.1) for training our models. Among the 1,000 in

Model trained on	DF				FSh				F2F				NT				FS				FF++ (Average)			
	One seg		Two seg		One seg		Two seg		One seg		Two seg		One seg		Two seg		One seg		Two seg		One seg		Two seg	
	IoU	AUC	IoU	AUC	IoU	AUC	IoU	AUC	IoU	AUC	IoU	AUC	IoU	AUC	IoU	AUC	IoU	AUC	IoU	AUC	IoU	AUC	IoU	AUC
DF (ours)	0.987	0.986	0.975	0.984	0.926	0.917	0.886	0.923	0.963	0.96	0.937	0.959	0.748	0.729	0.603	0.723	0.954	0.956	0.920	0.954	0.915	0.912	0.860	0.910
FSh (ours)	0.957	0.963	0.933	0.958	0.972	0.984	0.961	0.980	0.971	0.969	0.947	0.966	0.764	0.748	0.628	0.745	0.966	0.968	0.938	0.965	0.926	0.928	0.878	0.924
F2F (ours)	0.970	0.976	0.959	0.976	0.971	0.978	0.954	0.972	0.982	0.986	0.974	0.984	0.840	0.836	0.739	0.832	0.983	0.985	0.966	0.981	0.950	0.953	0.918	0.950
NT (ours)	0.971	0.985	0.961	0.980	0.969	0.983	0.958	0.979	0.963	0.980	0.955	0.977	0.949	0.974	0.933	0.966	0.946	0.970	0.932	0.965	0.960	0.979	0.949	0.974
FS (ours)	0.941	0.935	0.898	0.932	0.963	0.962	0.931	0.955	0.970	0.970	0.950	0.969	0.679	0.679	0.514	0.642	0.981	0.987	0.967	0.983	0.904	0.901	0.842	0.898
FF++ (CADDM)	0.966	0.987	0.957	0.987	0.943	0.981	0.903	0.977	0.935	0.971	0.891	0.971	0.745	0.883	0.598	0.874	0.974	0.986	0.970	0.986	0.913	0.962	0.864	0.869
FF++ (ours)	0.974	0.988	0.962	0.982	0.974	0.989	0.962	0.982	0.975	0.988	0.965	0.983	0.959	0.972	0.938	0.967	0.975	0.988	0.955	0.978	0.971	0.985	0.957	0.979

Table 2: Results for temporal segmentation on the proposed benchmark temporal deepfake dataset with **randomly chosen fake-segments**. Structure of this table is similar to of Table 1. For each test sub-dataset we have tested separately for videos with one fake-segment and two fake-segments from our proposed benchmark dataset (see section 3.2).

each sub-dataset within FF++, we have used 800 for training and validation, and we tested on the remaining 200. We have experimented with the videos with compression level ‘c23’ and all the frames were used in the training and testing, i.e., no frames were skipped.

To evaluate the temporal segmentation performance, we have used our proposed novel benchmark dataset (see Section 3.2) for testing the temporal segmentation performance. This dataset is generated based on the original FF++ dataset, where each video contains both real and fake segments. There are 5 sub-datasets within the temporal dataset, the same as the original FF++, each having 100 videos with an average of 633.9 frames per video, 24.3% fake frames for videos with one fake segment, and 41.1% fake frames for videos with two fake segments. We also experimented with other popular datasets: FF++, CelebDF[37], DFDC[18] and WildDeepFakes[74]. These datasets were used to compare our method’s performance in terms of traditional deepfake detection experiments: classical (same-dataset) deepfake detection and generalizability (cross-dataset).

For face extraction and alignment, we have used DLIB[29], and aligned face images were resized to 224×224 for all frames in the train and test sets.

Settings. Vision Transformer (ViT) and Shift and Scaling (SSF). A simple set of preprocessing steps including extraction of the frames and cropping the face region is done as data preparation for ViT. We use the ViT-B/16[20] architecture as the backbone for our ViT model and trained it using four A100 GPUs with a batch size of 64 for 80 epochs. The optimization algorithm used was ‘AdamW’ with a weight decay of 0.05. A warm-up strategy was applied for the learning rate, starting with $1e^{-7}$ for the first five epochs and then increasing linearly to $1e^{-3}$. The dropout probability was set to 0.1, and the input image size was set to 224×224 . To improve performance, an exponential moving average was used with a decay

rate of 0.99992 together with automatic mixed precision. The model was initialized with pre-trained weights on Imagenet-21K-SSF.

Timeseries settings. The dimension of the feature vector for each frame from the ViT was 768. After windowing these vectors as shown in Figure 2 the input dimension for the Timeseries Transformer (TsT) was $(W, 768)$. In our experiments, we used $W = 5$ however it is also possible to use different values for W . There are a total of 8 transformer blocks in the TsT, with an 8-headed attention connection. Each attention-head’s dimension is 512. After the transformer blocks, we have a one-dimensional global average pool prior to the MLP head. We have used batch size of 64, ‘categorical cross-entropy’ as the loss and ‘Adam’ as the optimizer with $1e^{-4}$ learning rate for training the TsT. We use an early stopping technique with patience of 10 to speed up the training procedure.

4.2. Temporal segmentation analysis

We have used our proposed benchmark dataset to test our method for the temporal segmentation problem, where we try to classify deepfake videos at the frame-level instead of video-level. The metrics we use to measure the performance for temporal segmentation are Intersection over Union (IoU) and Area under the ROC Curve (AUC). The baseline IoU (for random guessing of the class of a frame) is 1/3 as shown in Section 3.2.2. We have trained six separate models on six training sets: FaceForensics++ (FF++) and the five sub-datasets within FF++ i.e. Deepfakes (DF), Face-Shifter (FSh), Face2Face (F2F), Neural Textures (NT) and FaceSwap (FS). Similarly, we report the results for each model on the six test sets (FF++ and its five sub-datasets) in Tables 1 and 2. We compare our results with the CADDM [19] where we have used the original implementation and model weights published by the authors with a slight modification to compute frame-level AUC and IoU values.

	DF	FSh	F2F	NT	FS	FF++	C-DF	DFDC	WDF
DF	0.993	0.965	0.975	0.830	0.967	0.946	0.590	0.558	0.631
FSh	0.980	<i>0.990</i>	0.978	0.848	0.980	0.955	0.609	0.535	0.619
F2F	<i>0.990</i>	0.993	0.990	0.917	<i>0.993</i>	0.977	0.670	0.603	0.676
NT	0.973	0.967	0.967	<i>0.965</i>	0.960	0.967	<i>0.715</i>	<i>0.626</i>	<i>0.693</i>
FS	0.960	0.975	0.982	0.770	0.995	0.936	0.584	0.512	0.611
FF++	0.985	0.983	<i>0.983</i>	0.967	0.983	0.982	0.790	0.667	0.703

Table 3: Results (in AUC) for video level classification. Similar to Tables 1 and 2, each row represents models trained on specific training data and the columns constitute the test data. Along with FF++ and its sub-datasets we have tested each model on other datasets such as CelebDF (C-DF), DFDC, and WildDeepFakes (WDF).

As seen, each model does very well when it was tested on the test-set from the same dataset as it was trained on, hence the results on the diagonals are either the best or the second-best in every column while the second-best results are only lower in the range from 0.001 to 0.009. As expected, the model trained on the whole FF++ is the overall best-performing model. In Table 1 and 2 we can see that our detection method outperforms the latest state-of-the-art, CADDM [19] method in the temporal segmentation task.

However, the results of the other five models give us some important findings. Of the five sub-datasets in FF++, three were made with a face swapping technique (DF, FSh and FS) and the other two were made with face reenactment (F2F and NT). We can see that the models trained in F2F and NT perform better than the other three models. Since face reenactment deepfakes are devoid of strong artifacts compared to face swapping deepfakes, models trained on face reenactment methods tend to generalize well to other methods. Similarly, models trained on the face swapping methods do not generalize well to face re-enactment test data as seen on the ‘F2F’ and ‘NT’ columns in both Table 1 and 2. Overall, out of the five subdatasets of FF++, NT is the best one to train a detection model if others are unavailable.

Also, our method does significantly better with these two sub-datasets compared to CADDM [19] which indicates the necessity to learn temporal features to predict the transition between real and fake segments more effectively.

4.3. Video-level classification and Generalizability

The classical approach to deepfake detection has always been to predict the class (real or fake) of a deepfake video i.e. to make video-level predictions. We performed experiments on the test data from the original datasets and reported the results in Table 3 using AUC as the metric. We also measure our models’ performance on test data from datasets outside of FF++: CelebDF (C-DF), DFDC,

Method	CelebDF	FF++	NT
Xception [10]	0.653	0.997	0.842
SRM [40]	0.659	0.969	<i>0.943</i>
SPSL [39]	0.724	0.969	0.805
MADD [70]	0.674	0.998	–
SLADD [8]	0.797	<i>0.984</i>	–
CADDM [19]	0.931	0.998	0.837
Ours (ViT+TsT)	0.790	0.982	0.967

Table 4: Comparison with other state-of-the-art methods in terms of video-level AUC. Models were trained on FF++ and evaluated on CelebDF, FF++ and NT. Results for the other methods were taken from their own paper or github repository page. NT is the most challenging deepfake generation technique in terms of both temporal segmentation and video level detection. Our method significantly outperforms recent methods in detecting NT fakes while performing competitively on other datasets.

Length (seconds)	1.0	2.0	3.0	4.0	5.0	6.0	7.0	8.0	9.0	10.0
IoU	0.961	0.962	0.962	0.963	0.963	0.965	0.965	0.967	0.967	0.969
AUC	0.963	0.976	0.979	0.982	0.982	0.984	0.984	0.985	0.984	0.985

Table 5: IoU and AUC of the our proposed method across different lengths of deepfake segments. Our approach is largely robust to variation in the length of injected deepfake segment.

and WildDeepFakes (WDF). Results for the sub-datasets of FF++ (DF, FSh, F2F, NT and FS) follow the results of the temporal analysis where the diagonal values are the best or the second-best in a column, i.e. models when tested on data from the same sub-dataset generally perform very well. And, similar to the previous results (i.e. temporal segmentation), we see that models trained on face re-enactment data (F2F and NT) perform better than other sub-dataset-models when tested on unseen data, i.e. these two models generalize well in comparison with other models. However, we see the best results from tests on CelebDF, DFDC, and WildDeepFakes from the model trained on the full FF++ dataset.

It is also noticeable that the AUC scores for DFDC and WildDeepFakes are lower compared to CelebDF. While CelebDF is a dataset with videos made solely by the face swapping technique, DFDC is a combination of multiple methods such as face swapping (Deepfake Autoencoder and Morphable-mask), Neural talking-heads and GAN-based methods. WildDeepFakes dataset contains videos from the Internet which may contain videos generated using a variety of methods. Some of these methods are totally unseen due to their absence in the FF++ dataset.

We further evaluate and compare our results with the latest state-of-the-art methods in Table 4 using the model trained on FF++ and tested on FF++, CelebDF, and NT. Our

detection method outperforms the state-of-the-art methods in detecting videos generated by the NT method which is the most difficult method in FF++. Our method also performs very competitively with the most recent state-of-the-art deepfake detection methods [19, 8, 70] and outperforms most methods in video-level predictions in both same-dataset and cross-dataset scenarios. Our method comprehensively exceeds the performance of several state-of-the-art methods such as Xception [55], SRM [40], SPSL [39], MADD [70] and others (not included in Table 4) [48, 35, 41] in generalizability (i.e. cross-dataset/CelebDF). These results demonstrate that our method can also be used with high confidence for traditional deepfake detection and for unseen data (i.e. video-level) alongside temporal segmentation despite not being optimized for this objective.

4.4. Varying Lengths of Fake Segments

Our proposed method is effective in identifying even short segments of deepfake that can significantly alter the message conveyed by a video. To evaluate the performance of our method, we conducted experiments on a test set comprising 100 videos with varying lengths of fake segments, and the results are presented in Table 5 and Figure 5. Specifically, we create fake segments with durations ranging from 0.2 seconds to 19 seconds, with an increase of 0.2 seconds, and calculate the average IoU and AUC over the 100 videos. The increment of 0.2 seconds is the average duration of two phonemes in English [22], which we assume to be the unit duration for a fake segment. In particular, we did not use the smoothing of noisy frames in this experiment.

Our method achieves high accuracy in detecting very short fake-segments with a duration of less than 1.0 second, yielding an AUC value of over 0.91. Furthermore, as the length of the fake segment increases, our method performs even better in terms of AUC and IoU. This experiment provides evidence that our proposed method can identify even the slightest alterations in very short fake-segments, highlighting its effectiveness in detecting deepfake videos.

4.5. Ablation study

We conducted experiments to evaluate the effectiveness of our proposed method without TsT and the smoothing algorithm for both temporal segmentation detection and video-level detection. The model was trained on the full FF++ training data and tested on the proposed temporal segmentation benchmark dataset and FF++ test set for temporal segmentation and video-level detection, respectively. We used a MLP head on the ViT to classify frames for the experiment where the TsT was not included. Our ablated model achieved great results in both test sets, which are reported in Table 6. To provide a better comparison, we also reported the results from the full model with the TsT and smoothing algorithm. While we observe that the ViT al-

Model trained on	Temporal Evaluation (IoU)			Video level (AUC)	
	ViT	ViT+TsT	ViT+TsT+Smooth.	ViT	ViT+TsT
DF	0.960	0.967 (+0.007)	0.974 (+0.014)	0.973	0.985 (+0.012)
FSH	0.960	0.967 (+0.007)	0.974 (+0.014)	0.973	0.983 (+0.010)
F2F	0.956	0.968 (+0.012)	0.975 (+0.019)	0.973	0.983 (+0.010)
NT	0.933	0.948 (+0.015)	0.959 (+0.026)	0.965	0.967 (+0.002)
FS	0.951	0.965 (+0.014)	0.975 (+0.024)	0.973	0.983 (+0.010)
FF++	0.952	0.964 (+0.012)	0.971 (+0.019)	0.971	0.982 (+0.011)

Table 6: Ablation study on temporal segmentation of deepfakes and video-level classification. For temporal evaluation, results (IoU) from three experiments are reported: from Vision Transformer (ViT) only, with Time-series Transformer (TsT) and also including smoothing. Video level evaluation (in AUC) is reported for ViT, and ViT with TsT. Changes in the results are reported in bold and are in parentheses.

ready performs very well, a significant improvement can be seen in both temporal and video-level performance with the inclusion of the TsT and the smoothing algorithm.

Another ablation study is on varying window sizes for the TsT. In TsT we use a sliding window technique to accumulated features of multiple frames within a window so that the TsT can learn temporal features. We have experimented with varying window sizes in terms of number of frames, accompanying with varying overlap values in the sliding window method. Based on the results from these experiments on both frame-level and video-level predictions as reported in Tables 11 and 12 we have selected the window size of 5 frames with overlap of 4 frames to be the optimal parameters.

5. Discussion

While most methods tackle deepfake detection at the video-level, we propose a robust and generalizable method that can produce results at the frame, segment and entire video level. This allows maximal flexibility in analyzing content for the presence of deepfakes and additionally provides comparison points for future research along these related but separate evaluation protocols.

Our method is based on supervised pretraining of the image encoder, which limits computational requirements in two forms. First, the image encoder is trained independently on individual video frames with frame-level supervision, nullifying the need to learn temporal relationships between frames. Second, a large part of the backbone is frozen and initialized using readily available weights from ImageNet, considerably reducing the computational cost of obtaining a deepfake related representation in the encoder. The proposed method achieves robust IoU metrics across the proposed single and multi-segment deepfakes, while maintaining competitive performance on video-level deepfake detection and generalized deepfake detection.

References

- [1] Darius Afchar, Vincent Nozick, Junichi Yamagishi, and Isao Echizen. Mesonet: a compact facial video forgery detection network. In *2018 IEEE international workshop on information forensics and security (WIFS)*, pages 1–7. IEEE, 2018. 2
- [2] Saadaldeen Rashid Ahmed, Emrullah Sonuç, Mohammed Rashid Ahmed, and Adil Deniz Duru. Analysis survey on deepfake detection and recognition with convolutional neural networks. In *2022 International Congress on Human-Computer Interaction, Optimization and Robotic Applications (HORA)*, pages 1–7. IEEE, 2022. 2
- [3] Irene Amerini, Leonardo Galteri, Roberto Caldelli, and Alberto Del Bimbo. Deepfake video detection through optical flow based cnn. In *Proceedings of the IEEE/CVF international conference on computer vision workshops*, 2019. 2
- [4] Roberto Amoroso, Davide Morelli, Marcella Cornia, Lorenzo Baraldi, Alberto Del Bimbo, and Rita Cucchiara. Parents and children: Distinguishing multimodal deepfakes from natural images. *arXiv preprint arXiv:2304.00500*, 2023. 2
- [5] Belhassen Bayar and Matthew C Stamm. A deep learning approach to universal image manipulation detection using a new convolutional layer. In *Proceedings of the 4th ACM workshop on information hiding and multimedia security*, pages 5–10, 2016. 2
- [6] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023. 2
- [7] Zhixi Cai, Kalin Stefanov, Abhinav Dhall, and Munawar Hayat. Do you really mean that? content driven audio-visual deepfake dataset and multimodal method for temporal forgery localization. In *2022 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–10. IEEE, 2022. 3
- [8] 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, 7, 8
- [9] Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. Adaptformer: Adapting vision transformers for scalable visual recognition. *arXiv preprint arXiv:2205.13535*, 2022. 3
- [10] 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
- [11] Komal Chugh, Parul Gupta, Abhinav Dhall, and Ramathanan Subramanian. Not made for each other-audio-visual dissonance-based deepfake detection and localization. In *Proceedings of the 28th ACM international conference on multimedia*, pages 439–447, 2020. 3
- [12] Umur Aybars Ciftci, Ilke Demir, and Lijun Yin. Fakecatcher: Detection of synthetic portrait videos using biological signals. *IEEE transactions on pattern analysis and machine intelligence*, 2020. 2
- [13] Davide Cozzolino, Giovanni Poggi, and Luisa Verdoliva. Recasting residual-based local descriptors as convolutional neural networks: an application to image forgery detection. In *Proceedings of the 5th ACM workshop on information hiding and multimedia security*, pages 159–164, 2017. 2
- [14] Davide Cozzolino, Andreas Rössler, Justus Thies, Matthias Nießner, and Luisa Verdoliva. Id-reveal: Identity-aware deepfake video detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15108–15117, 2021. 2
- [15] FaceSwap-GAN developer. Faceswap-gan. <https://github.com/shaoanlu/faceswap-GAN>. Accessed: 2023-02-16. 2
- [16] DFaker Developers. Dfaker. <https://github.com/dfaker/df>. Accessed: 2023-02-16. 2
- [17] Deepfakes Faceswap Developers. Deepfakes. <https://github.com/deepfakes/faceswap>. Accessed: 2023-02-16. 2, 4
- [18] Brian Dolhansky, Joanna Bitton, Ben Pfau, Jikuo Lu, Russ Howes, Menglin Wang, and Cristian Canton Ferrer. The deepfake detection challenge (dfdc) dataset. *arXiv preprint arXiv:2006.07397*, 2020. 6
- [19] Shichao Dong, Jin Wang, Renhe Ji, Jiajun Liang, Haoqiang Fan, and Zheng Ge. Towards a robust deepfake detector: Common artifact deepfake detection model. *arXiv preprint arXiv:2210.14457*, 2022. 5, 6, 7, 8
- [20] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 6
- [21] Ricard Durall, Margret Keuper, Franz-Josef Pfreundt, and Janis Keuper. Unmasking deepfakes with simple features. *arXiv preprint arXiv:1911.00686*, 2019. 2
- [22] Gunnar Fant, Anita Kruckenberg, and Lennart Nord. Durational correlates of stress in swedish, french and english. *Journal of phonetics*, 19(3-4):351–365, 1991. 8
- [23] Weinan Guan, Wei Wang, Jing Dong, Bo Peng, and Tieniu Tan. Collaborative feature learning for fine-grained facial forgery detection and segmentation. *arXiv preprint arXiv:2304.08078*, 2023. 2
- [24] Yinan He, Bei Gan, Siyu Chen, Yichun Zhou, Guojun Yin, Luchuan Song, Lu Sheng, Jing Shao, and Ziwei Liu. Forgerynet: A versatile benchmark for comprehensive forgery analysis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4360–4369, 2021. 3
- [25] Javier Hernandez-Ortega, Ruben Tolosana, Julian Fierrez, and Aythami Morales. Deepfakes detection based on heart rate estimation: Single-and multi-frame. In *Handbook of Digital Face Manipulation and Detection: From DeepFakes to Morphing Attacks*, pages 255–273. Springer International Publishing Cham, 2022. 3
- [26] Anubhav Jain, Nasir Memon, and Julian Togelius. A dataless faceswap detection approach using synthetic images. In *2022*

- IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–7. IEEE, 2022. 2
- [27] 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
- [28] Minha Kim, Shahroz Tariq, and Simon S Woo. Fretal: Generalizing deepfake detection using knowledge distillation and representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1001–1012, 2021. 2
- [29] Davis E King. Dlib-ml: A machine learning toolkit. *The Journal of Machine Learning Research*, 10:1755–1758, 2009. 6
- [30] Pavel Korshunov and Sébastien Marcel. Deepfakes: a new threat to face recognition? assessment and detection. *arXiv preprint arXiv:1812.08685*, 2018. 2
- [31] Pavel Korshunov and Sébastien Marcel. Improving generalization of deepfake detection with data farming and few-shot learning. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 4(3):386–397, 2022. 2
- [32] Marek Kowalski. Faceswap. <https://github.com/MarekKowalski/FaceSwap>. Accessed: 2023-02-16. 4
- [33] 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. 2, 4
- [34] 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
- [35] Xiaodan Li, Yining Lang, Yuefeng Chen, Xiaofeng Mao, Yuan He, Shuhui Wang, Hui Xue, and Quan Lu. Sharp multiple instance learning for deepfake video detection. In *Proceedings of the 28th ACM international conference on multimedia*, pages 1864–1872, 2020. 8
- [36] Yuezun Li, Ming-Ching Chang, and Siwei Lyu. In icu oculi: Exposing ai created fake videos by detecting eye blinking. In *2018 IEEE international workshop on information forensics and security (WIFS)*, pages 1–7. IEEE, 2018. 2
- [37] 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. 6
- [38] Dongze Lian, Daquan Zhou, Jiashi Feng, and Xinchao Wang. Scaling & shifting your features: A new baseline for efficient model tuning. *arXiv preprint arXiv:2210.08823*, 2022. 1, 3
- [39] 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*, pages 772–781, 2021. 7, 8
- [40] 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. 7, 8
- [41] 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. 8
- [42] Falko Matern, Christian Riess, and Marc Stamminger. Exploiting visual artifacts to expose deepfakes and face manipulations. In *2019 IEEE Winter Applications of Computer Vision Workshops (WACVW)*, pages 83–92. IEEE, 2019. 2
- [43] Aakash Varma Nadimpalli and Ajita Rattani. On improving cross-dataset generalization of deepfake detectors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 91–99, 2022. 2
- [44] Huy H Nguyen, Junichi Yamagishi, and Isao Echizen. Use of a capsule network to detect fake images and videos. *arXiv preprint arXiv:1910.12467*, 2019. 2
- [45] Yotam Nitzan, Kfir Aberman, Qirui He, Orly Liba, Michal Yarom, Yossi Gandelsman, Inbar Mosseri, Yael Pritch, and Daniel Cohen-Or. Mystyle: A personalized generative prior. *ACM Transactions on Graphics (TOG)*, 41(6):1–10, 2022. 2
- [46] Leandro A Passos, Danilo Jodas, Kelton AP da Costa, Luis A Souza Júnior, Danilo Colombo, and João Paulo Papa. A review of deep learning-based approaches for deepfake content detection. *arXiv preprint arXiv:2202.06095*, 2022. 2
- [47] Ivan Perov, Daiheng Gao, Nikolay Chervoniy, Kunlin Liu, Sugasa Marangonda, Chris Umé, Mr Dpfks, Carl Shift Facenheim, Luis RP, Jian Jiang, et al. Deepfacelab: Integrated, flexible and extensible face-swapping framework. *arXiv preprint arXiv:2005.05535*, 2020. 2, 4
- [48] Yuyang Qian, Guojun Yin, Lu Sheng, Zixuan Chen, and Jing Shao. Thinking in frequency: Face forgery detection by mining frequency-aware clues. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII*, pages 86–103. Springer, 2020. 8
- [49] Md Atiqur Rahman and Yang Wang. Optimizing intersection-over-union in deep neural networks for image segmentation. In *International symposium on visual computing*, pages 234–244. Springer, 2016. 5
- [50] Nicolas Rahmouni, Vincent Nozick, Junichi Yamagishi, and Isao Echizen. Distinguishing computer graphics from natural images using convolution neural networks. In *2017 IEEE workshop on information forensics and security (WIFS)*, pages 1–6. IEEE, 2017. 2
- [51] Md Shohel Rana, Mohammad Nur Nobi, Beddhu Murali, and Andrew H Sung. Deepfake detection: A systematic literature review. *IEEE access*, 10:25494–25513, 2022. 2
- [52] Christian Rathgeb, Ruben Tolosana, Ruben Vera-Rodriguez, and Christoph Busch. *Handbook of digital face manipulation and detection: from DeepFakes to morphing attacks*. Springer Nature, 2022. 2
- [53] Hamid Rezaatofghi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding

- box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 658–666, 2019. 5
- [54] 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*, pages 10684–10695, 2022. 2
- [55] Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. FaceForensics++: Learning to detect manipulated facial images. In *International Conference on Computer Vision (ICCV)*, 2019. 3, 4, 5, 8
- [56] Sanjay Saha and Terence Sim. Is face recognition safe from realizable attacks? In *2020 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–8, 2020. 2
- [57] Mahmood Sharif, Sruti Bhagavatula, Lujio Bauer, and Michael K Reiter. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proceedings of the 2016 acm sigsac conference on computer and communications security*, pages 1528–1540, 2016. 2
- [58] 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
- [59] Laura Stroebel, Mark Llewellyn, Tricia Hartley, Tsui Shan Ip, and Mohiuddin Ahmed. A systematic literature review on the effectiveness of deepfake detection techniques. *Journal of Cyber Security Technology*, 7(2):83–113, 2023. 2
- [60] Zekun Sun, Yujie Han, Zeyu Hua, Na Ruan, and Weijia Jia. Improving the efficiency and robustness of deepfakes detection through precise geometric features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3609–3618, 2021. 2
- [61] 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. 2, 4
- [62] 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. 4
- [63] 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. 2, 4
- [64] Rotem Tzaban, Ron Mokady, Rinon Gal, Amit Bermanto, and Daniel Cohen-Or. Stitch it in time: Gan-based facial editing of real videos. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022. 2
- [65] 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. 3
- [66] Luisa Verdoliva. Media forensics and deepfakes: an overview. *IEEE Journal of Selected Topics in Signal Processing*, 14(5):910–932, 2020. 2
- [67] Ying Xu, Kiran Raja, Luisa Verdoliva, and Marius Pedersen. Learning pairwise interaction for generalizable deepfake detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 672–682, 2023. 2
- [68] Xin Yang, Yuezun Li, and Siwei Lyu. Exposing deep fakes using inconsistent head poses. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8261–8265. IEEE, 2019. 2
- [69] Peipeng Yu, Zhihua Xia, Jianwei Fei, and Yujiang Lu. A survey on deepfake video detection. *Iet Biometrics*, 10(6):607–624, 2021. 2
- [70] Hanqing Zhao, Wenbo Zhou, Dongdong Chen, Tianyi Wei, Weiming Zhang, and Nenghai Yu. Multi-attentional deepfake detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2185–2194, 2021. 2, 7, 8
- [71] Tianchen Zhao, Xiang Xu, Mingze Xu, Hui Ding, Yuanjun Xiong, and Wei Xia. Learning self-consistency for deepfake detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 15023–15033, 2021. 2
- [72] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks, 2017. 2
- [73] Wanyi Zhuang, Qi Chu, Zhentao Tan, Qiankun Liu, Haojie Yuan, Changtao Miao, Zixiang Luo, and Nenghai Yu. Uia-vit: Unsupervised inconsistency-aware method based on vision transformer for face forgery detection. In *European Conference on Computer Vision*, pages 391–407. Springer, 2022. 2
- [74] Bojia Zi, Minghao Chang, Jingjing Chen, Xingjun Ma, and Yu-Gang Jiang. Wilddeepfake: A challenging real-world dataset for deepfake detection. In *Proceedings of the 28th ACM international conference on multimedia*, pages 2382–2390, 2020. 6