

TALL: Thumbnail Layout for Deepfake Video Detection

Yuting Xu^{1,3*}, Jian Liang^{2,4}, Gengyun Jia⁵, Ziming Yang^{1,3}, Yanhao Zhang⁶, Ran He^{2,4†}

¹ Institute of Information Engineering, Chinese Academy of Sciences

² CRIPAC & MAIS, Institute of Automation, Chinese Academy of Sciences

³ School of Cyber Security, UCAS ⁴ School of Artificial Intelligence, UCAS

⁵ School of Communications and Information Engineering, NJUPT ⁶ OPPO Research Institute

yuting.xu@cripac.ia.ac.cn, liangjian92@gmail.com, rhe@nlpr.ia.ac.cn

Abstract

The growing threats of deepfakes to society and cybersecurity have raised enormous public concerns, and increasing efforts have been devoted to this critical topic of deepfake video detection. Existing video methods achieve good performance but are computationally intensive. This paper introduces a simple yet effective strategy named Thumbnail Layout (TALL), which transforms a video clip into a pre-defined layout to realize the preservation of spatial and temporal dependencies. Specifically, consecutive frames are masked in a fixed position in each frame to improve generalization, then resized to sub-images and rearranged into a pre-defined layout as the thumbnail. TALL is model-agnostic and extremely simple by only modifying a few lines of code. Inspired by the success of vision transformers, we incorporate TALL into Swin Transformer, forming an efficient and effective method TALL-Swin. Extensive experiments on intra-dataset and cross-dataset validate the validity and superiority of TALL and SOTA TALL-Swin. TALL-Swin achieves 90.79% AUC on the challenging cross-dataset task, FaceForensics++ \rightarrow Celeb-DF. The code is available at <https://github.com/rainy-xu/TALL4Deepfake>.

1. Introduction

Deepfakes generate and manipulate facial appearances to deceive viewers through generation techniques [61, 48]. With the remarkable success of generative adversarial networks [14, 27], deepfake products have become photo-realistic that humans can not distinguish. These deepfake products [21, 49] may be misused for malicious purposes, leading to severe trust issues and security problems, such as financial fraud, identity theft, and celebrity impersonation [53, 40].




Figure 1. **The AUC and FLOPs trade-off of different backbones.** Image-level backbones with TALL enjoy comparable accuracy-cost trade-offs with the 3DCNN and video transformer family on the unseen Celeb-DF dataset. All models with the same setting are trained on the FF++ (HQ) dataset.

ation [53, 40]. The rapid development of social media exacerbates the abuse of deepfakes. Therefore, it is crucial to develop advanced detection methods to protect the data privacy of individual users.

Most previous image-based methods [22, 69] perform well on intra-dataset, but their generalizability needs to be improved. Recent research has focused on video-based methods to detect deepfake by modeling spatio-temporal dependencies. There are subtle spatio-temporal inconsistencies between frames since the deepfake algorithms are executed frame by frame. The core of video-level approaches for deepfake detection is capturing inconsistencies through temporal modeling. Existing deepfake video detection methods generally follow two directions. Some methods [15, 16] use two-branch networks or modules to learn spatial and temporal information separately and then fuse them. However, these two-branch approaches may frag-

*This work was done when she was a student in CRIPAC.

†Corresponding author.

ment spatiotemporal cooperation and lead to subtle artifacts being neglected. Others directly use classic temporal models such as LSTM and 3D-CNNs. These methods are computationally intensive. The current rise of transformers for vision task backbones has prompted the emergence of corresponding deepfake detection methods. They are accompanied by significant computational complexity that makes them challenging to deploy and use, despite breakthroughs in performance. To enjoy benefits from both image and video methods, we are curious to see whether it is possible to append information about the temporal dimension to the image dimension.

This work develops a simple yet effective Thumbnail Layout (TALL) for deepfake detection by spatio-temporal modeling. TALL is computationally cheap and retains both temporal and spatial information. In detail, we use dense sampling to extract multiple clips in the video and then randomly select four consecutive frames in the video segment. Subsequently, a block is masked at a fixed position in each frame. Finally, the frames are resized as sub-image and sequentially rearranged into a pre-defined layout as a thumbnail, which has the same size as the clip frames. As shown in Figure 1, TALL brings two advantages compared to the previous spatio-temporal modeling methods for deepfake detection: (1) TALL contains local and global contextual deepfake patterns. (2) TALL is a model-agnostic method for spatio-temporal modeling deepfake patterns at zero computation and zero parameters.

Furthermore, we discover that the better temporal modeling capabilities backbone has, the better performance TALL achieves. Based on the proposed TALL, we complement a baseline for video deepfake detection based on Swin Transformer [36], called TALL-Swin. We validate TALL-Swin on four popular benchmark datasets, including FaceForensics++, Celeb-DF, DFDC, and DeeperForensics. Our method gains a remarkable improvement over the state-of-the-art approaches. The main contributions of our paper are summarized as follows:

- We provide a new perspective for an efficient strategy for video deepfake detection called Thumbnail Layout (TALL), which incorporates both spatial-temporal dependencies, and allows the model to capture spatial-temporal inconsistencies.
- We propose a spatio-temporal modeling method called TALL-Swin, which efficiently captures the inconsistencies between deepfake video frames.
- Extensive experiments demonstrate the validity of our proposed TALL and TALL-Swin. TALL-Swin outperforms previous methods in both intra-dataset and cross-dataset scenarios.

2. Related Work

2.1. Image-Level Deepfake Detection

Typically, existing deepfake detection methods fall into two categories: image-level and video-level methods. The image-level methods [25, 13] always exploit the artifacts of deepfake images in the spatial domain, such as discrepancies between local regions [42, 59], grid-like structure in frequency space [10], and differences in global texture statistics [37] that provide specific clues to distinguish deepfakes from the real images. F3Net [43] and FDFL [31] utilize the same pipeline that utilizes frequency-aware features and RGB information to capture the traces in different input spaces separately. RFM [54] and Multi-att [65] propose an attention-guided data augmentation mechanism to guide detectors to discover undetectable deepfake clues. Face X-ray [32] and PCL [67] provide effective ways to outline the boundary of the forged face for detecting deepfakes. ICT [11] exploits an identity extraction module to detect identity inconsistency in the suspect image. Similarly, M2tr [55] detects local inconsistencies within frames at different spatial levels. Generally, image-level methods suffer over-fitting issues when a specific technique manipulates the images, and they ignore temporal information.

2.2. Video-Level Deepfake Detection

To improve the generalization of deepfake detectors, many studies generate diversity and generic deepfake data, while other studies capture the temporal incoherence of fake videos as generic clues. Some recent works propose detecting temporal inconsistency using well-designed spatio-temporal neural networks, and others [15, 16] attempt to add modules to image models that capture temporal information. STIL [15] formulates deepfake video detection as a spatial and temporal inconsistency learning process and integrates both spatial and temporal features in a unified 2D CNN framework. FTCN [68] detects temporal-related artifacts instead of spatial artifacts to promote generalization. LipForensics [18] is proposed to learn high-level semantic irregularity in mouth movement in the generated video. RealForensics [17] uses auxiliary data sets during training in exchange for generalization at the cost of higher computational demands. The video-based methods achieve strong generalization but suffer from large computational overhead. To reduce computational costs, we propose TALL which gathers consecutive video frames into thumbnails for learning spatio-temporal consistency.

2.3. Deepfake Detection with Vision Transformer

Recently, ViT [12] has achieved impressive performance in computer vision tasks [24, 23, 44]. Many studies extend the ViT for deepfake detection [66, 58]. These methods achieve better performance compared to CNN-based

models, but also sacrifice computational efficiency. A few works [55, 66] attempt to extend the transformer for deepfake detection due to the advent of the visual transformer (ViT) and the impressive ability to model long-range data, different from two-branch architectures that capture short-range and long-range temporal inconsistencies with a single-branch model. ICT [11] aims to detect identity consistency in deepfake video but may fail in detecting face reenactment and entire face synthesis results. DFLL [28] extract the UV texture map to help the transformer to detect deepfakes, which may disrupt the continuity between video frames. DFTD [29] leverages ViT to consider both global and local information but ignores the problem of excessive model arithmetic requirements. Although the transformer-based approaches achieve promising performance, they are accompanied by significant computational complexity that makes them challenging to deploy and use, and the long-range dependencies may be insufficiently exploited in detection models. Swin Transformer [36] produces a hierarchical feature representation and has linear computational complexity concerning input image size, which is suitable as a general-purpose backbone for various vision tasks. In this paper, we cooperate with Swin Transformer to form our robust and efficient method TALL-Swin.

3. Method

TALL is a deepfake video detection strategy that transforms a video clip into an all-in-one thumbnail without the extra computational overhead. In the following sections, we begin with the motivation of TALL for deepfake detection in Section 3.1. Then we present the technical details of the TALL in Section 3.2. Finally, a generalizable Swin-TALL baseline is introduced to explore subtle artifacts in Section 3.3. More details of the layout design are presented in the Appendix.

3.1. Motivation

While recent studies have attempted to address noticeable flaws through techniques like slight motion blurring and temporal consistency loss, subtle spatio-temporal artifacts still remain. These artifacts are important for detecting deepfakes, but they introduce two problems: 1) video-based models are less efficient, and 2) analyzing information over long distances may overlook local artifacts, which are critical for deepfake detection. To address these challenges, we propose the TALL strategy, which naturally incorporates temporal information into image-level tasks without disrupting spatial information. This approach enables the image-level model to detect deepfakes in videos. Furthermore, we discovered that TALL provides even greater performance gains when combined with a powerful spatial model, resulting in the TALL-Swin.

In detail, TALL arranges consecutive frames in the

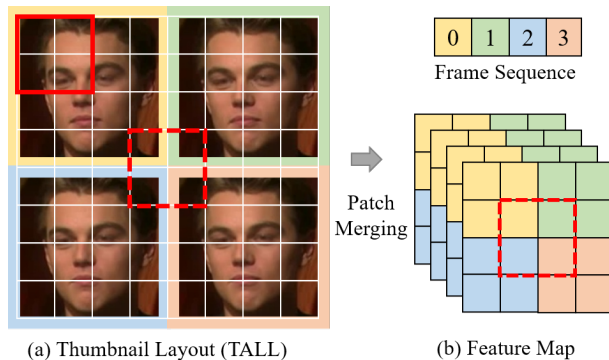


Figure 2. Illustration of the TALL and shifted window process for computing self-attention in the TALL.

temporal order in a compact 2×2 layout, in line with the calculation theory of convolution and shifted window. TALL contains both spatial and temporal information so that model can learn both intra-frame artifacts and inter-frame inconsistency and obtains comparable performance to video-based methods. Here we use the shifted window to explain TALL’s mechanism. As illustrated in Figure 2 (a), the model computes self-attention while accounting for spatial dependencies across sub-images (represented by the solid red box). When the window spans multiple sub-images (represented by the red dash box), the model is able to capture temporal inconsistencies between frames. Moreover, TALL leverages both local and global contexts of deepfake patterns to ensure robust modeling capabilities for short and long-range spatial dependencies. Compared to previous methods, we anticipate that TALL strikes a balance between speed and accuracy, sacrificing a little spatial information while preserving performance. Based on the fact that attention-based models are better at handling contextual features and that the Swin-Transformer uses shifted windows to reduce computation and memory, we further complement TALL-Swin baseline for video deepfake detection.

3.2. Thumbnail Layout (TALL)

Given a video $V \in \mathbb{R}^{T \times C \times H \times W}$, where T is the frame length of the video, C is the number of channels, and $H \times W$ is the resolution of the frames. Assuming each video contains N clips, we divide a video into N equal segments of length T/N and then sample consecutive t (set to 4 by default) frames from the segments at random locations to form one clip. Then, the thumbnail I is rearranged of sub-images ($C \times \frac{H}{\sqrt{t}} \times \frac{W}{\sqrt{t}}$) that are resized from the above t frames. To maximize the utility of TALL, we mask the organized N square masks of the thumbnail. It is based on two core designs: 1) The position of the masks is random between different sub-images, which retains the advantages of the Cutout [8] that encourages the network to focus more

Algorithm 1 Pseudocode of TALL in a PyTorch-like style.

```
# x: one clip of video (T*C*H*W)
# T: frame number of clip
# C: channels; s: mask size
# d: length of frame included in the
  thumbnail
# r: rows of thumbnail
# x_tall: thumbnail image (224*224)

#TALL's augmentation strategy
h = np.random.randint(H)
w = np.random.randint(W)
#the mask position is fixed for each frame
m = np.ones((H, W))
h1 = np.clip(h - s // 2, 0, H)
h2 = np.clip(h + s // 2, 0, H)
w1 = np.clip(w - s // 2, 0, W)
w2 = np.clip(w + s // 2, 0, W)
m[h1:h2, w1:w2] = 0
m = torch.from_numpy(m)
m = mask.expand_as(x)
x = x * m
#TALL: generation of the thumbnail
x = x.view(-1, H, W).unsqueeze(0)
x = x.view((-1, C*d) + x_tall.size()[2:])
x = rearrange(x, 'b (th tw c) h w
  -> b c (th h) (tw w)', th=r, c=C)
x_tall = interpolate(x, size=H)
```

on complementary and less prominent features. 2) We fix the position of the mask within a clip to take advantage of the fact that most deepfake videos are frame-by-frame tampered with, thus forcing the model to detect inconsistencies between adjacent frames of the deepfake videos. We do not allow the mask to appear on the seams of the thumbnail but allow for partial mask inclusion in the thumbnail. The detailed procedure of TALL is summarized in Algorithm 1.

3.3. TALL-Swin

To balance efficiency and model performance for spatio-temporal feature learning and to leverage the benefits of attention-based models, we enhanced a baseline deepfake detection model called TALL-Swin by incorporating the Swin Transformer [36]. Given the characteristics of TALL, we slightly modified the window size of Swin-B in TALL-Swin. We first enlarge the window size of the first three stages of the model so that the interaction between frames in the thumbnail becomes more frequent, forcing the model to learn more detailed spatio-temporal dependencies. Next, we set the window size of the last stage to be the same as the feature map size, enabling the window to perform global attention computations while TALL-Swin captures global spatial-temporal dependencies. As a result, the size of the last layer of the feature map became smaller, reducing the window size without introducing any additional computational overhead. Consequently, the window sizes for the four stages of TALL-Swin are [14, 14, 14, 7]. Note that the patch merging process makes TALL-Swin captures a more

comprehensive range of dependencies through hierarchical representations, as shown in Figure 2 (b).

Given a video of length T , each frame contains N patches, and the window contains P patches. To demonstrate the superiority of TALL-Swin in terms of computational consumption, we show below the computational complexity of the image-level transformer and video-level transformer, including ViT [12], Swin [36], ViViT [2], and TALL-Swin respectively:

$$\begin{aligned}\Omega_{\text{ViT}} &= 4TNC^2 + 2TN^2C, \\ \Omega_{\text{Swin}} &= 4TNC^2 + 2TPNC, \\ \Omega_{\text{ViViT}} &= 4TNC + 2T^2N^2C, \\ \Omega_{\text{TALL-Swin}} &= TNC^2 + \frac{1}{2}TPNC.\end{aligned}\tag{1}$$

TALL-Swin has the lowest computational complexity compared to image and video-level transformer methods. Subsequent experiments will demonstrate that TALL-Swin maintains performance, albeit at the sacrifice of some spatial information.

The cross-entropy loss is employed to optimize the TALL-Swin, which is defined as:

$$\mathcal{L}_{CE} = -\frac{1}{n} \sum_{i=1}^n y_i \log \mathcal{F}(x_i) + (1 - y_i) (\log(1 - \mathcal{F}(x_i))), \tag{2}$$

where x_i indicates input clip, y_i denotes the label of clip, n is the number of clip, \mathcal{F} is TALL-Swin.

4. Experiments

4.1. Setup

Datasets. Following previous works [18, 17, 68], we evaluate the TALL and TALL-Swin on four widely used datasets. **FaceForensics++** [45] is a most-used benchmark on intra-dataset deepfake detection, consisting of 1,000 real videos and 4,000 fake videos in four different manipulations: DeepFake [7], FaceSwap [38], Face2Face [52], and NeuralTextures [51]. Besides, FaceForensics++ contains multiple video qualities, *e.g.* high quality (HQ), low quality (LQ) and RAW. **Celeb-DF (CDF)** [34] is a popular benchmark on cross-dataset, which contains 5,693 deepfake videos generated from celebrities. The improved compositing process was used to improve the various visual artifacts presented in the video. Celeb-DF is also suitable for deepfake detection tasks with a reference set. **DFDC** [9] is a large-scale benchmark developed for Deepfake Detection Challenge. This dataset includes 124k videos from 3,426 paid actors. The existing deepfake detection methods do perform not very well on DFDC due to their sophisticated deepfake techniques. **DeeperForensics (DFo)** [26] includes 60,000 videos with 17.6 million frames for deepfake detection, whose videos vary in identity, pose, expression, emotion, lighting conditions, and blend shape with high quality.

Models	Temp.	CDF	DFDC	FLOPs	Params	PT
I3D-RGB* [4]	✓	78.24	65.58	222.7G	25M	1K
R3D-50* [19]	✓	79.63	67.73	296.6G	46M	1K
ResNet50* [20]	×	76.38	64.01	25.5G	21M	1K
+TALL	✓	80.90	65.54	25.5G	21M	1K
EffNetB4* [50]	×	78.19	66.81	8.3G	19M	1K
+TALL	✓	83.37	67.15	8.3G	19M	1K
VTN [41]	✓	83.20	73.50	296.6G	46M	21K
VidTR [63]	✓	83.30	73.30	117G	93M	21K
ViViT* [2]	✓	86.96	74.61	628G	310M	21K
ISTVT [64]	✓	84.10	74.20	455.8G	-	-
ViT-B* [12]	×	82.33	72.64	55.4G	84M	21K
+TALL	✓	86.58	74.10	55.4G	84M	21K
Swin-B* [36]	×	83.13	73.01	47.5G	86M	21K
TALL-Swin	✓	90.79	76.78	47.5G	86M	21K

Table 1. **Performance of different backbones.** TALL consistently improves the accuracy over different image-level models. We show the AUC, FLOPs, and number of parameters for each model on the cross-dataset scenario. All models are trained on FF++ (HQ). ✓ indicates the model enables temporal modeling. * indicates our implementation. PT indicates pre-train. 1K and 21K indicate the model pre-trained on ImageNet-1K and 21K respectively. The best results are **bold**.

Implementation Details. We use MTCNN to detect face for each frame in the deepfake videos, only extract the maximum area bounding box and add 30% face crop size from each side as in LipForensics [18]. The ImageNet-21K pretrained Swin-B model is used as our backbone. Excluding ablation experiments, we sample 8 clips using dense sampling, each clip contains 4 frames. The size of the thumbnail is 224×224 . Following Swin Transformer [36], Adam [30] optimization is used with a learning rate of $1.5e-5$ and batch size of 4, using a cosine decay learning rate scheduler and 10 epochs of linear warm-up. We adopt Acc. (accuracy) and AUC (Area Under Receiver Operating Characteristic Curve) as the evaluation metrics for extensive experiments. To ensure a fair comparison, we calculate video-level predictions for the image-based method and average the predictions across the entire video (following previous works [18, 16, 35, 68]). Note that results are directly cited from published papers if we follow the same setting.

4.2. Scaling over Backbones

To verify our assumption, we adopt several image-level backbones commonly used for deepfake detection for comparison with the video-level backbones. As shown in Table 1 above the double horizontal line, we first compare the accuracy and complexity of the CNN-based video and image backbones. Although I3D [4] and R3D [19] achieve

better performance than vanilla ResNet50 [20] and EfficientNet [50], the computation costs are huge, such as R3D-50 with 296G FLOPs. For ResNet and EfficientNet who added TALL, ResNet achieves better AUC both on CDF (76.38 VS 80.93) and DFDC (64.01 VS 65.54) datasets, and EfficientNet achieves 5.18% better AUC on CDF.

The second section contains the video and image transformers. Compared to video transformers, the image-based ViT and Swin fail to achieve better performance due to the lack of temporal modeling. For example, ViViT achieves 86.96% AUC on CDF, which is 3.6% higher than Swin although ViViT with $13 \times$ more computation. By way of contrast, ViT+TALL achieves 86.58% AUC on CDF with 55.4G FLOPs, which is comparable to AUC with ViViT but with low computation. Accordingly, Swin’s performance was significantly improved with the addition of TALL without computation increment. On the other hand, TALL boosts higher performance on models with learned long-range dependencies. *e.g.*, ResNet+TALL (+4.5% on CDF and +1.5% on DFDC) vs. Swin+TALL (+7.6% on CDF and +3.6% on DFDC). These two section results demonstrate that TALL provides both spatial and temporal information and enables the model to learn spatial and temporal inconsistencies for video deepfake detection.

Methods	FF++ (HQ)		FF++ (LQ)	
	Acc.	AUC	Acc.	AUC
MesoNet [1]	83.10	-	70.47	-
Xception [6]	95.73	96.30	86.86	89.30
Face X-ray [32]	-	87.35	-	61.60
Two-branch [39]	96.43	98.70	86.34	86.59
Add-Net [70]	96.78	97.74	87.50	91.01
F3-Net [43]	97.52	98.10	90.43	90.43
FDL [31]	96.69	99.30	89.00	92.40
Multi-Att [65]	97.60	99.29	88.69	90.40
RECCE [3]	97.06	99.32	91.03	95.02
LipForensics [18]	98.80	99.70	94.20	98.10
DFDT [29]	98.18	99.26	92.67	94.48
ADT [56]	92.05	96.30	81.48	82.52
ST-M2TR [55]	-	99.42	-	95.31
VTN [41]	98.47	-	94.02	-
VidTR [63]	97.42	-	92.12	-
ViViT* [2]	92.60	-	88.02	-
ISTVT [64]	99.00	-	96.15	-
TALL-Swin	98.65	99.87	92.82	94.57

Table 2. **Intra-dataset evaluations.** We report the video-level Acc. (%) and AUC (%) on the FF++ dataset. HQ indicates high quality, and LQ indicates low quality.

4.3. Comparison with State-of-the-art Methods

Intra-dataset evaluations. Following ISTVT [64], we show the results of the FF++ dataset under both Low Quality (LQ) and High Quality (HQ) videos, and report comparisons against several advanced methods in Table 2. We can observe that advanced video-based transformers have better results than CNN-based methods. Compared to video-based transformer methods, TALL-Swin has comparable performance and lower consumption to the previous video transformer method with HQ settings. However, TALL-Swin gets unsatisfactory results with the LQ setting. The LQ setting is obtained by severely compressing the videos. So the reason for the result may be that TALL scales the frame to a smaller size, causing more spatial information to be lost in the frame. We will investigate the possibility of other designs to further improve performance in the LQ setting.

Method	CDF	DFDC	FSh	DFo	Avg.
Xception [6]	73.70	70.90	72.00	84.50	75.28
CNN-aug [57]	75.60	72.10	65.70	74.40	71.95
CNN-GRU [46]	69.80	68.90	80.80	74.10	73.40
Patch-based [5]	69.60	65.60	57.80	81.80	68.70
Face X-Ray [32]	79.50	65.50	92.80	86.80	81.15
Multi-Att [65]	75.70	68.10	66.00	77.70	71.88
DSP-FWA [33]	69.50	67.30	65.50	50.20	63.13
LipForensics [18]	82.40	73.50	97.10	97.60	87.65
<hr/>					
FTCN [68]	86.90	74.00	98.80	98.80	89.63
RealForensics [17]	86.90	75.90	99.70	99.30	90.45
DFDT [29]	88.30	76.10	97.80	96.90	89.70
<hr/>					
VTN [41]	83.20	73.50	98.70	97.70	88.30
VidTR [63]	83.50	73.30	98.00	97.90	88.10
ViViT* [2]	86.96	74.61	99.41	99.19	90.05
ISTVT [64]	84.10	74.20	99.30	98.60	89.10
TALL-Swin	90.79	76.78	99.67	99.62	91.71

Table 3. **Generalization to unseen datasets.** We report the video-level AUC (%) on four unseen datasets: Celeb-DF (CDF), DFDC, FaceShifter (FSh), and DeeperForensics (DFo).

Generalization to unseen datasets. In addition to the intra-dataset comparisons, we also investigate the generalization ability of our method. Adhering to the deepfake video detection cross-dataset protocol [18], we train a model on FF++ (HQ) then test on Celeb-DF (CDF), DFDC, FaceShifter (FSh), and DeeperForensics (DFo) datasets. As shown in Table 3: (1) Video-based methods generally have better results than image-based methods, which shows that temporal information is helpful for the deepfake video detection task. For example, Lip outperforms Face X-ray’s AUC by a wide margin. In addition, most transformer-based models have higher performance than CNN-based models. For the transformer-based models, both achieved

an average AUC of 88%, while the best CNN-based video-level models only achieved 87%. (2) TALL-Swin achieves state-of-the-art results on Celeb-DF, DFDC, and DeeperForensics datasets, and also beats its competitors on Celeb-DF dataset by a large margin (3.8%). The results demonstrate that TALL-Swin performs well when encountering unseen datasets with better generalization ability than previous video transformer methods.

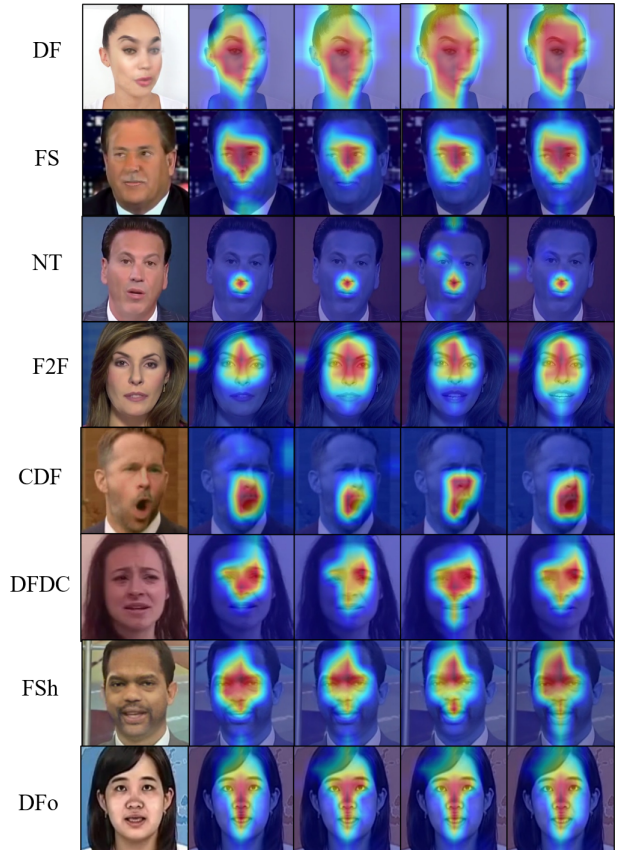


Figure 3. **Saliency map visualization of TALL-Swin on different datasets.** The first four rows of samples are from the FF++ dataset, and the last four rows are from the unseen datasets.

Analysis of saliency map visualization. We adopt Grad-CAM [47] to visualize where the TALL-Swin is paying its attention to the deepfake faces. In Figure 3, we give the results on intra-dataset and cross-dataset scenarios. All models are trained on FF++ (HQ). It can be observed in the first four rows of Figure 3 that TALL-Swin captures method-specific artifacts. Note that the DF transfers the face region from a source video to a target, and the NT only modifies the facial expressions corresponding to the mouth region. TALL-Swin corresponds to focus on the face region and the mouth region. Furthermore, our model traces the more generalized artifacts that are independent of manipulation methods, *e.g.*, blending boundaries (CDF), and abnormal motions in the clip (DFDC, FSh, Dfo).

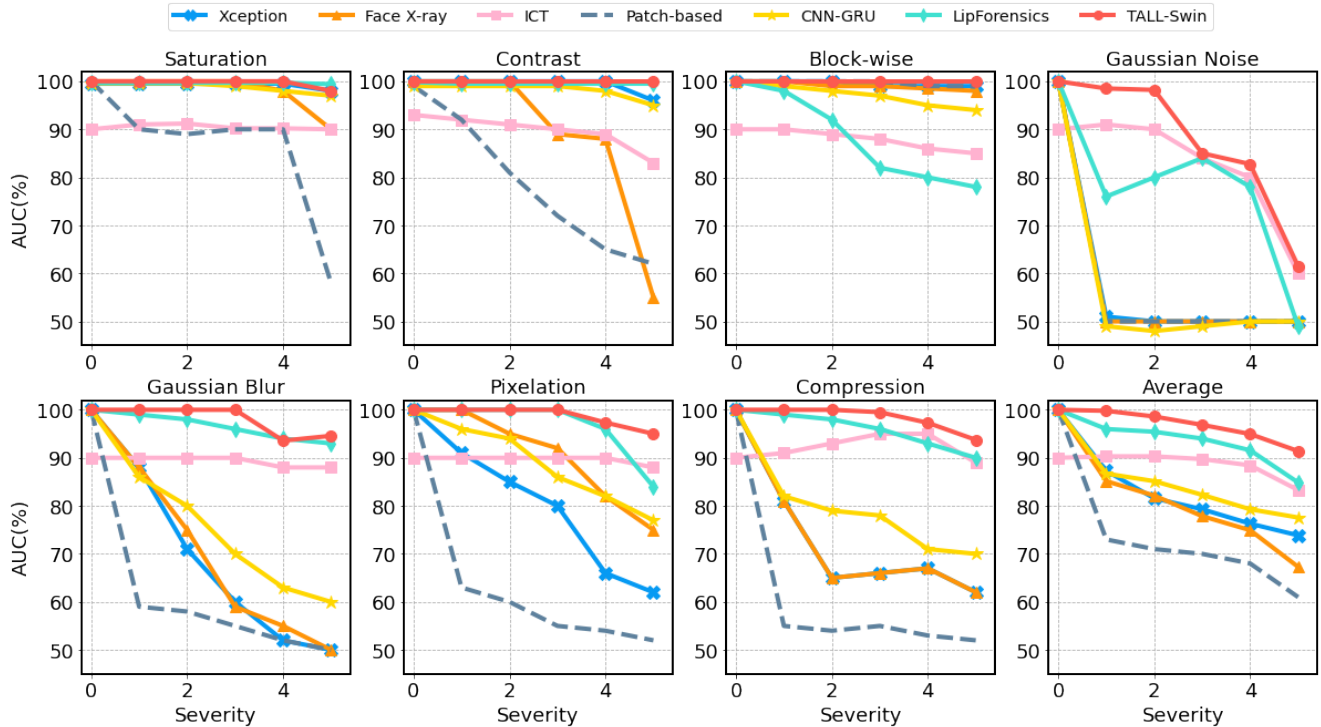


Figure 4. **Robustness to various unseen corruptions.** We report the video-level AUC (%) of our methods under five different levels of seven particular types of corruption. “Average” denotes the mean across all corruptions at each severity level. Our TALL-Swin is more robust than previous methods for all corruptions.

Method	Clean	Saturation	Contrast	Block	Noise	Blur	Pixel	Compress	Avg.
Xception (ICCV’19) [6]	99.8	99.3	98.6	99.7	53.8	60.2	74.2	62.1	78.3
CNN-GRU (CVPRW’19) [46]	99.9	99.0	98.8	97.9	47.9	71.5	86.5	74.5	82.3
CNN-aug (CVPR’20) [57]	99.8	99.3	99.1	95.2	54.7	76.5	91.2	72.5	84.1
Patch-based (ECCV’20) [5]	99.9	84.3	74.2	99.2	50.0	54.4	56.7	53.4	67.5
Face X-ray (CVPR’20) [32]	99.8	97.6	88.5	99.1	49.8	63.8	88.6	55.2	77.5
LipForensics (ICCV’21) [18]	99.9	99.9	99.6	87.4	73.8	96.1	95.6	95.6	92.5
FTCN (ICCV’21) [68]	99.4	99.4	96.7	97.1	53.1	95.8	98.2	86.4	89.5
RealForensics (CVPR’22) [17]	99.8	99.8	99.6	98.9	79.7	95.3	98.4	97.6	95.6
TALL-Swin w/o mask	100.0	100.0	100.0	99.8	83.5	97.3	98.4	97.9	96.7
TALL-Swin	100.0	100.0	100.0	100.0	85.3	97.6	98.5	98.1	97.1

Table 4. **Average robustness to unseen corruptions.** Average Video-level AUC (%) across five intensity levels for each corruption type proposed in DFo [26]. “Avg” indicates the mean across all corruptions and all levels.

Robustness to unseen perturbations. Deepfake detectors must be robust to common perturbations, given that video propagation on social media causes video compression, noise addition, etc. We also study the performance of robustness to unseen perturbations. Following RealForensics [17], the experiment applies seven unseen perturbations to fake videos at five intensity levels. In Figure 4, we show results of increasing the severity of each corruption. We can observe that other methods degrade dramatically as the perturbations become more severe. TALL-Swin still has a high performance. However, TALL-Swin degrades when

the Gaussian noise reaches level five. Table 4 presents the average AUC across all intensity levels for corruption types. We observe that our method is significantly more robust to most perturbations than other methods. The good robustness may be from both the design of TALL and the proposed mask augmentation. The main reason may be the consecutive multi-frame input. We empirically consider that the key to deepfake detection is local inconsistency, the continuous frame design has less redundant information, ensuring that the model finds locally important clues.




Figure 5. Illustration of different layout designs.

4.4. Ablation Study

We perform the ablation study to analyze the effects of each component and hyper-parameter in TALL-Swin. All experiments are trained on FF++ (HQ) and tested on the CDF and DFDC datasets.

Effects of different layouts. We train a TALL-Swin model on FF++(HQ) for each layout illustrated in Figure 5, to analyze in which layout of the thumbnails the model learns the strongest generalization of the spatial-temporal dependence of the deepfake patterns. As shown in Table 5, the model with a compact layout like Figure 5 (d) has good generalization ability on the unseen datasets. A compact layout like Figure 5 (d) may help the model to learn the temporal dependence across frames because such a form provides the shortest distance between any two images.

Layout	CDF	DFDC
Figure 5 (a)	85.52	70.02
Figure 5 (b)	84.93	73.57
Figure 5 (c)	86.66	72.12
Figure 5 (d)	87.60	74.32

Table 5. Effects of different layouts. All models here are trained without mask augmentation.

Study on the numbers of sub-image. We use Swin-B as the baseline for this study to compare the effect of different thumbnail layout schemes on the model’s generalization ability. Changing frames to thumbnails involves scaling, so we also investigate the impact of resizing and random cropping pre-processing on model performance. We set up four variants: resizing pre-process with 4×4 layout, 3×3 layout and 2×2 layout; random cropping pre-process with 2×2 layout. As shown in Table 6, the model performance degrades sharply when using 4×4 layout. This may be due to the small size of each sub-image that the spatial information is not captured well by the model. The result of 3×3 layout also slightly decreases. 2×2 layout with resizing pre-processing beats 2×2 layout with random crop. We also found that TALL-Swin achieves the best performance and the AUC score increases 3.2% compared to the baseline, suggesting that thumbnails in a 2×2 layout are more helpful to TALL-Swin than original frames.

Effects of Sub-image’s size. We eliminate the scaling

Pre-process	Layout	CDF	DFDC
None	-	83.13	73.01
Resize	4×4	80.18	70.45
Resize	3×3	83.18	72.98
Crop	2×2	78.55	73.30
Resize	2×2	87.60	74.32

Table 6. Study on the numbers of sub-image. All models here are trained without mask augmentation.

operation for sub-images to allow for more flexible layout settings. However, we’ve observed that when the number of sub-images grows at their original size, the performance improvements are only slight. Additionally, the computational complexity increases dramatically with the number of frames (4.3 times more than the TALL setting), as demonstrated in Table 7. To strike a balance between performance and computational complexity, we reduce the resolution of sub-images in TALL.

Subimage-size	Layout	FLOPs	CDF	DFDC
224×224	3×3	253G	88.69	75.98
224×224	2×2	185G	88.15	75.01
112×112	2×2	47.5G	87.60	74.32

Table 7. Effects of Sub-image’s size. All models here are trained without mask augmentation.

Study on absence and order of thumbnails. In this case, we study the impact of missing the last sub-image and the last two sub-images on the model’s performance. The first two rows of Table 8 show that all four sub-images contribute to the model performance. Besides, we set the order of the different thumbnails to evaluate the TALL-Swin. We consider three orders: forward, reverse, and random. Forward order performs the best for three different orders. This may be because of the positional encoding of different frames in TALL-Swin.

Variants	CDF	DFDC
0, 1, 2, -	86.46	69.51
0, 1, -, -	84.22	69.09
Random	85.85	70.30
Reverse	86.65	72.37
Forward	87.60	74.32

Table 8. Ablation study of absence and order of thumbnails. All models here are trained without mask augmentation.

Effects of different orders on other backbones. In order to prove that the phenomenon is not incidental, we also conduct experiments on ResNet50 and EfficientNet for three different orders as shown in Table 9. As expected, for-

ward order outperforms reverse and random orders both on ResNet50 and EfficientNet, which indicates that TALL can learn the temporal dependency.

Variants	CDF	DFDC
ResNet50+TALL	76.38	64.01
Random	78.14	64.12
Reverse	78.54	64.87
Forward	80.90	65.54
EfficientNet+TALL	78.19	66.81
Random	81.01	66.13
Reverse	81.66	66.69
Forward	83.37	67.15

Table 9. Ablation studies of different orders of TALL.

Effectiveness of mask strategy. In this work, TALL-Swin is trained on the FF++ (HQ) dataset without any data enhancement as the baseline except for Multi-scale Crop and Random Horizontal Flip. To validate the effectiveness of the mask strategy, we compare our default baseline with different data augmentation strategies: 1) The Cutout [8] on one sub-image; 2) The Cutout on four sub-images. 3) The combination of Mixup [62] and Cutmix [60] on four sub-images, as shown in Table 10. The performance of a random Cutout [8] on four sub-images is better than on one sub-image. Besides, the mask strategy leads to better performance than the well-known Cutout (1.46%). This supports our hypothesis that strategy encourages models to learn subtle temporal-spatial variations and improves model generalization ability. Further, our augmentation strategy exceeds 1.02% than the combination of Mixup and Cutmix, demonstrating the augmentation’s effectiveness in TALL for video detection.

Augmentation	Count	CDF	DFDC
None	-	87.60	74.32
Cutout	1	89.06	74.07
Cutout	4	89.33	75.22
Mixup+CutMix	4	89.75	75.33
TALL’s mask	4	90.79	76.78

Table 10. Study of the augmentation strategy in TALL. The count column represents the number of blocks on the thumbnail.

Study on window size. We study the effect of window size on model performance and computational cost. The results are shown in Table 11. Our window expansion for the first three phases will increase the model performance by 1.74% AUC. The results in the second and third rows show that the first three stages of the window getting the largest would not give a boost to the model. Our analysis of a too-large window may lead to a weakening of the model’s ability to learn local information in the sub-image.

Window size	CDF	DFDC
(7,7,7,7)	85.60	73.32
(14,14,14,7)	87.60	74.32
(28,28,28,7)	86.65	74.21

Table 11. Ablation studies of window size. All models here are trained without mask augmentation.

5. Conclusion

This paper presents a novel perspective on detecting deepfake videos using TALL. TALL is both simple and effective, enabling joint spatio-temporal modeling without any additional costs. TALL representation reveals normal deepfake patterns with local-global contextual features. We further propose a new baseline for deepfake video detection called TALL-Swin, which efficiently captures the inconsistencies between deepfake video frames. Extensive experiments demonstrate that TALL-Swin achieves promising results for various unseen deepfake types and strong robustness to a wide range of common corruptions.

Acknowledgment

This work was partially funded by National Natural Science Foundation of China under Grants (62276256, U21B2045 and U20A20223) and Beijing Nova Program under Grant Z211100002121108. The authors wish to thank Huaibo Huang and Lijun Sheng in no particular order, for insightful discussions.

References

- [1] Darius Afchar, Vincent Nozick, Junichi Yamagishi, and Isao Echizen. Mesonet: a compact facial video forgery detection network. In *Proc. WIFS*, pages 1–7, 2018. 5
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lucic, and Cordelia Schmid. Vivit: A video vision transformer. In *Proc. ICCV*, pages 6836–6846, 2021. 4, 5, 6
- [3] Junyi Cao, Chao Ma, Taiping Yao, Shen Chen, Shouhong Ding, and Xiaokang Yang. End-to-end reconstruction-classification learning for face forgery detection. In *Proc. CVPR*, pages 4113–4122, 2022. 5
- [4] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proc. CVPR*, pages 6299–6308, 2017. 5
- [5] Lucy Chai, David Bau, Ser-Nam Lim, and Phillip Isola. What makes fake images detectable? understanding properties that generalize. In *Proc. ECCV*, pages 103–120, 2020. 6, 7
- [6] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proc. CVPR*, pages 1251–1258, 2017. 5, 6, 7
- [7] deepfakes. Deepfakes. <https://github.com/deepfakes/faceswap>. 2021-11-13. 4

- [8] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv:1708.04552*, 2017. 3, 9
- [9] Brian Dolhansky, Joanna Bitton, Ben Pflaum, Jikuo Lu, Russ Howes, Menglin Wang, and Cristian Canton Ferrer. The deepfake detection challenge (dfdc) dataset. *arXiv:2006.07397*, 2020. 4
- [10] Chengdong Dong, Ajay Kumar, and Eryun Liu. Think twice before detecting gan-generated fake images from their spectral domain imprints. In *Proc. CVPR*, pages 7865–7874, 2022. 2
- [11] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Ting Zhang, Weiming Zhang, Nenghai Yu, Dong Chen, Fang Wen, and Baining Guo. Protecting celebrities from deepfake with identity consistency transformer. In *Proc. CVPR*, pages 9468–9478, 2022. 2, 3
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proc. ICLR*, 2021. 2, 4, 5
- [13] Jianwei Fei, Yunshu Dai, Peipeng Yu, Tianrun Shen, Zhihua Xia, and Jian Weng. Learning second order local anomaly for general face forgery detection. In *Proc. CVPR*, pages 20270–20280, 2022. 2
- [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. In *Proc. NeurIPS*, volume 27, 2014. 1
- [15] Zhihao Gu, Yang Chen, Taiping Yao, Shouhong Ding, Jilin Li, Feiyue Huang, and Lizhuang Ma. Spatiotemporal inconsistency learning for deepfake video detection. In *Proc. ACM MM*, pages 3473–3481, 2021. 1, 2
- [16] Zhihao Gu, Yang Chen, Taiping Yao, Shouhong Ding, Jilin Li, and Lizhuang Ma. Delving into the local: Dynamic inconsistency learning for deepfake video detection. In *Proc. AAI*, pages 744–752, 2022. 1, 2, 5
- [17] Alexandros Haliassos, Rodrigo Mira, Stavros Petridis, and Maja Pantic. Leveraging real talking faces via self-supervision for robust forgery detection. In *Proc. CVPR*, 2022. 2, 4, 6, 7
- [18] Alexandros Haliassos, Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Lips don’t lie: A generalisable and robust approach to face forgery detection. In *Proc. CVPR*, pages 5039–5049, 2021. 2, 4, 5, 6, 7
- [19] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Learning spatio-temporal features with 3d residual networks for action recognition. In *Proc. ICCV Workshops*, pages 3154–3160, 2017. 5
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. CVPR*, pages 770–778, 2016. 5
- [21] Gee-Sern Hsu, Chun-Hung Tsai, and Hung-Yi Wu. Dual-generator face reenactment. In *Proc. CVPR*, pages 642–650, 2022. 1
- [22] Yihao Huang, Felix Juefei-Xu, Qing Guo, Yang Liu, and Geguang Pu. Fakelocator: Robust localization of gan-based face manipulations. *IEEE Transactions on Information Forensics and Security*, 17:2657–2672, 2022. 1
- [23] Ge-Peng Ji, Guobao Xiao, Yu-Cheng Chou, Deng-Ping Fan, Kai Zhao, Geng Chen, and Luc Van Gool. Video polyp segmentation: A deep learning perspective. *MIR*, 19(6):531–549, 2022. 2
- [24] Ge-Peng Ji, Mingchen Zhuge, Dehong Gao, Deng-Ping Fan, Christos Sakaridis, and Luc Van Gool. Masked vision-language transformer in fashion. *MIR*, 20(3):421–434, 2023. 2
- [25] Gengyun Jia, Meisong Zheng, Chuanrui Hu, Xin Ma, Yuting Xu, Luoqi Liu, Yafeng Deng, and Ran He. Inconsistency-aware wavelet dual-branch network for face forgery detection. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3(3), 2021. 2
- [26] Liming Jiang, Ren Li, Wayne Wu, Chen Qian, and Chen Change Loy. Deeperforensics-1.0: A large-scale dataset for real-world face forgery detection. In *Proc. CVPR*, pages 2889–2898, 2020. 4, 7
- [27] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proc. CVPR*, pages 4401–4410, 2019. 1
- [28] Sohail Ahmed Khan and Hang Dai. Video transformer for deepfake detection with incremental learning. In *Proc. ACM MM*, pages 1821–1828, 2021. 3
- [29] Aminollah Khormali and Jiann-Shiun Yuan. Dfdt: An end-to-end deepfake detection framework using vision transformer. *Applied Sciences*, page 2953, 2022. 3, 5, 6
- [30] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014. 5
- [31] Jiaming Li, Hongtao Xie, Jiahong Li, Zhongyuan Wang, and Yongdong Zhang. Frequency-aware discriminative feature learning supervised by single-center loss for face forgery detection. In *Proc. CVPR*, pages 6458–6467, 2021. 2, 5
- [32] 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 *Proc. CVPR*, pages 5001–5010, 2020. 2, 5, 6, 7
- [33] Yuezun Li and Siwei Lyu. Exposing deepfake videos by detecting face warping artifacts. In *Proc. CVPR Workshops*, pages 656–663, 2019. 6
- [34] Yuezun Li, Xin Yang, Pu Sun, Honggang Qi, and Siwei Lyu. Celeb-df: A large-scale challenging dataset for deepfake forensics. In *Proc. CVPR*, pages 3207–3216, 2020. 4
- [35] 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 *Proc. CVPR*, pages 772–781, 2021. 5
- [36] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proc. ICCV*, pages 10012–10022, 2021. 2, 3, 4, 5
- [37] Zhengzhe Liu, Xiaojuan Qi, and Philip HS Torr. Global texture enhancement for fake face detection in the wild. In *Proc. CVPR*, pages 8060–8069, 2020. 2

- [38] MarekKowalski. Faceswap. <https://github.com/MarekKowalski/FaceSwap/>. 2021-11-13. 4
- [39] Iacopo Masi, Aditya Killekar, Royston Marian Mascarenhas, Shenoy Pratik Gurudatt, and Wael AbdAlmageed. Two-branch recurrent network for isolating deepfakes in videos. In *Proc. ECCV*, pages 667–684, 2020. 5
- [40] Yisroel Mirsky and Wenke Lee. The creation and detection of deepfakes: A survey. *ACM CSUR*, 54(1):1–41, 2021. 1
- [41] Daniel Neimark, Omri Bar, Maya Zohar, and Dotan Asselmann. Video transformer network. In *Proc. ICCV*, pages 3163–3172, 2021. 5, 6
- [42] Yuval Nirkin, Lior Wolf, Yosi Keller, and Tal Hassner. Deepfake detection based on discrepancies between faces and their context. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):6111–6121, 2021. 2
- [43] Yuyang Qian, Guojun Yin, Lu Sheng, Zixuan Chen, and Jing Shao. Thinking in frequency: Face forgery detection by mining frequency-aware clues. In *Proc. ECCV*, pages 86–103, 2020. 2, 5
- [44] Rameswar Panda Quanfu Fan, Richard Chen. Can an image classifier suffice for action recognition? In *ICLR*, 2022. 2
- [45] Andreas Rossler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Niessner. Faceforensics++: Learning to detect manipulated facial images. In *Proc. ICCV*, pages 1–11, 2019. 4
- [46] Ekraam Sabir, Jiaxin Cheng, Ayush Jaiswal, Wael AbdAlmageed, Iacopo Masi, and Prem Natarajan. Recurrent convolutional strategies for face manipulation detection in videos. In *Proc. CVPR Workshops*, pages 80–87, 2019. 6, 7
- [47] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proc. ICCV*, pages 618–626, 2017. 6
- [48] Yichun Shi, Xiao Yang, Yangyue Wan, and Xiaohui Shen. Semanticstylegan: Learning compositional generative priors for controllable image synthesis and editing. In *Proc. CVPR*, pages 11254–11264, 2022. 1
- [49] Jingxiang Sun, Xuan Wang, Yong Zhang, Xiaoyu Li, Qi Zhang, Yebin Liu, and Jue Wang. Fenerf: Face editing in neural radiance fields. In *Proc. CVPR*, pages 7672–7682, 2022. 1
- [50] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *Proc. ICML*, pages 6105–6114, 2019. 5
- [51] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM TOG*, 38(4):1–12, 2019. 4
- [52] Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Face2face: Real-time face capture and reenactment of rgb videos. In *Proc. ICCV*, pages 2387–2395, 2016. 4
- [53] Luisa Verdoliva. Media forensics and deepfakes: an overview. *IEEE Journal of Selected Topics in Signal Processing*, 14(5):910–932, 2020. 1
- [54] Chengrui Wang and Weihong Deng. Representative forgery mining for fake face detection. In *Proc. CVPR*, pages 14923–14932, 2021. 2
- [55] Junke Wang, Zuxuan Wu, Wenhao Ouyang, Xintong Han, Jingjing Chen, Yu-Gang Jiang, and Ser-Nam Li. M2tr: Multi-modal multi-scale transformers for deepfake detection. In *Proc. ICMR*, pages 615–623, 2022. 2, 3, 5
- [56] Ping Wang, Kunlin Liu, Wenbo Zhou, Hang Zhou, Honggu Liu, Weiming Zhang, and Nenghai Yu. Adt: Anti-deepfake transformer. In *Proc. ICASSP*, pages 2899–1903, 2022. 5
- [57] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. Cnn-generated images are surprisingly easy to spot... for now. In *Proc. CVPR*, 2020. 6, 7
- [58] Deressa Wodajo and Solomon Atnafu. Deepfake video detection using convolutional vision transformer. *arXiv:2102.11126*, 2021. 2
- [59] Ziming Yang, Jian Liang, Yuting Xu, Xiao-Yu Zhang, and Ran He. Masked relation learning for deepfake detection. *IEEE Transactions on Information Forensics and Security*, pages 1696–1708, 2023. 2
- [60] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proc. ICCV*, pages 6023–6032, 2019. 9
- [61] Yu Zeng, Zhe Lin, and Vishal M Patel. Sketchedit: Mask-free local image manipulation with partial sketches. In *Proc. CVPR*, pages 5951–5961, 2022. 1
- [62] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. Mixup: Beyond empirical risk minimization. *Proc. ICLR*, 2018. 9
- [63] Yanyi Zhang, Xinyu Li, Chunhui Liu, Bing Shuai, Yi Zhu, Biagio Brattoli, Hao Chen, Ivan Marsic, and Joseph Tighe. Vidtr: Video transformer without convolutions. In *Proc. ICCV*, pages 13577–13587, 2021. 5, 6
- [64] Cairong Zhao, Chutian Wang, Guosheng Hu, Haonan Chen, Chun Liu, and Jinhui Tang. Istvt: Interpretable spatial-temporal video transformer for deepfake detection. *IEEE Transactions on Information Forensics and Security*, 18:1335–1348, 2023. 5, 6
- [65] Hanqing Zhao, Wenbo Zhou, Dongdong Chen, Tianyi Wei, Weiming Zhang, and Nenghai Yu. Multi-attentional deepfake detection. In *Proc. CVPR*, pages 2185–2194, 2021. 2, 5, 6
- [66] Hanqing Zhao, Wenbo Zhou, Dongdong Chen, Weiming Zhang, and Nenghai Yu. Self-supervised transformer for deepfake detection. *arXiv:2203.01265*, 2022. 2, 3
- [67] Tianchen Zhao, Xiang Xu, Mingze Xu, Hui Ding, Yuanjun Xiong, and Wei Xia. Learning self-consistency for deepfake detection. In *Proc. CVPR*, pages 15023–15033, 2021. 2
- [68] Yinglin Zheng, Jianmin Bao, Dong Chen, Ming Zeng, and Fang Wen. Exploring temporal coherence for more general video face forgery detection. In *Proc. ICCV*, pages 15044–15054, 2021. 2, 4, 5, 6, 7
- [69] Xiangyu Zhu, Hao Wang, Hongyan Fei, Zhen Lei, and Stan Z Li. Face forgery detection by 3d decomposition. In *Proc. CVPR*, pages 2929–2939, 2021. 1
- [70] Bojia Zi, Minghao Chang, Jingjing Chen, Xingjun Ma, and Yu-Gang Jiang. Wilddeepfake: A challenging real-world dataset for deepfake detection. In *Proc. ACM MM*, pages 2382–2390, 2020. 5