

FakeRadar: Probing Forgery Outliers to Detect Unknown Deepfake Videos

Zhaolun Li¹ Jichang Li^{2*} Yinqi Cai³ Junye Chen³
Xiaonan Luo¹ Guanbin Li^{3,2,5} Rushi Lan^{1,4*}

¹Guilin University of Electronic Technology ²Pengcheng Laboratory ³Sun Yat-sen University ⁴Guangxi Key Laboratory of Image and Graphic Intelligent Processing ⁵Guangdong Key Laboratory of Big Data Analysis and Processing

zhaolunli@mails.guet.edu.cn, li.jichang@pcl.ac.cn, rslan2016@163.com

Abstract

In this paper, we propose FakeRadar, a novel deepfake video detection framework designed to address the challenges of cross-domain generalization in real-world scenarios. Existing detection methods typically rely on manipulationspecific cues, performing well on known forgery types but exhibiting severe limitations against emerging manipulation techniques. This poor generalization stems from their inability to adapt effectively to unseen forgery patterns. To overcome this, we leverage large-scale pretrained models (e.g. CLIP) to proactively probe the feature space, explicitly highlighting distributional gaps between real videos, known forgeries, and unseen manipulations. Specifically, FakeRadar introduces Forgery Outlier Probing, which employs dynamic subcluster modeling and cluster-conditional outlier generation to synthesize outlier samples near boundaries of estimated subclusters, simulating novel forgery artifacts beyond known manipulation types. Additionally, we design Outlier-Guided Tri-Training, which optimizes the detector to distinguish real, fake, and outlier samples using proposed outlier-driven contrastive learning and outlierconditioned cross-entropy losses. Experiments show that FakeRadar outperforms existing methods across various benchmark datasets for deepfake video detection, particularly in cross-domain evaluations, by handling the variety of emerging manipulation techniques.

1. Introduction

Recently, generative techniques [28, 58, 62] for creating deepfake videos have advanced rapidly, producing highly realistic facial forgeries that challenge human perceptual limits [8, 23]. However, the malicious use of such technology, such as fabricated news and identity fraud, poses a significant threat to digital media credibility [7, 9]. Ear-

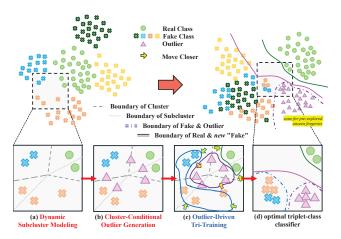


Figure 1. Conceptual pipeline of FakeRadar. (a)-(b) Forgery Outlier Probing first performs dynamic subcluster modeling on these features to partition finer-grained subclusters, then generates cluster-conditional outliers for each estimated subcluster to simulate unseen forgeries, emulating unknown manipulations. (c)-(d) The model is optimized using outlier-driven tri-training, allowing the detector to establish optimal boundaries for the real, fake, and outlier classes, with the latter two merged into a new "Fake" class during testing.

lier deepfake detection methods primarily rely on the passive learning of specific forgery artifacts, such as boundary inconsistencies [27, 54], blinking anomalies [32], and texture irregularities [17, 35]. While these approaches perform well against conventional deepfake manipulations, including Face2Face [46] and NeuralTextures [47], they exhibit critical vulnerabilities when confronted with emerging generative paradigms, such as diffusion models [55]. For instance, boundary-based detection methods [30, 44] fail to identify faces generated by advanced GANs (e.g., StyleGAN [25]). Similarly, blink analysis-based techniques [32] are rendered ineffective when processing temporally smoothed fake videos [57]. These limitations stem from a fundamental shortcoming: existing methods learn a constrained decision boundary in the feature space, where

^{*}Corresponding Authors.

known forgery traces occupy only a localized region, while novel manipulation techniques introduce artifacts in arbitrary locations. As a result, they struggle to generalize to new manipulation techniques that introduce artifacts in previously unseen ways, leading to significant performance degradation in cross-domain scenarios [11, 12].

To better understand its underlying cause, [2, 36] have analyzed the feature distributions of real and synthetic data. These studies reveal a fundamental embedding-space disparity: real data exhibit compact, dense distributions, whereas deepfake forgeries form scattered, sparse clusters due to their generation with diverse manipulation types, as demonstrated in Figure 1. This difference in distribution poses a significant challenge for existing detection methods, which are typically optimized to classify previously observed forgeries. Since deepfake forgeries do not adhere to a fixed pattern but rather introduce an infinite variety of modifications, classifiers trained on known forgery patterns struggle to recognize novel deepfakes that deviate from past observations. Therefore, a robust detection approach must move beyond merely learning pre-defined forgery patterns and instead develop the capability to identify anomalies that fall outside of previously known distributions. Given the limitations of existing methods in handling unseen forgery types, we propose a paradigm shift in deepfake detection. Instead of passively fitting to known forgery patterns, our approach proactively explores potential anomaly regions in the feature space. Inspired by radar systems that scan for unknown targets through frequency-spectrum probing, our method enhances detection robustness against novel deepfake manipulations.

In this work, we propose a new deepfake video detection framework, called FakeRadar, which leverages the deep feature priors of large-scale pre-trained models (i.e. CLIP [40]) to proactively "probe" unknown forgeries and enhance cross-domain generalization. Unlike conventional detectors [2, 5, 51] that perform only binary classification, FakeRadar explicitly introduces "forgery outliers" during training to simulate potential unknown forgeries and thus "pre-explore and cover" the latent space. As shown in Figure 1 (a) and (b), we first employ Forgery Outlier Probing, which conducts dynamic subcluster modeling on the features of real and known deepfake videos. By approximating the embedding space via fine-grained Gaussian mixture distributions, this algorithm automatically partitions training samples into distinct subclusters. Next, through clusterconditional outlier generation, Forgery Outlier Probing synthesizes outlier samples near the boundaries of these estimated subclusters with a certain "novel offset", thereby covering a broader range of potential forgery regions.

Additionally, FakeRadar adopts a model optimization strategy called Outlier-Guided Tri-Training, enabling the model to discriminate among "Real", "Fake", and "Out-

lier" samples, with its details as illustrated in Figure 1 (c) and (d). Specifically, Outlier-Guided Tri-Training proposes an outlier-driven contrastive loss and an outlier-conditioned cross-entropy loss for end-to-end model optimization. The former loss enhances the distance between real, fake, and outlier features via contrastive learning, emphasizing their differences, while the latter one ensures that the model has clear decision boundaries for the three categories, especially ensuring that outliers are not incorrectly classified as real. In this triplet-class (Real/Fake/Outlier) setting, the model can independently label unknown forgeries during training. During inference, both "Fake" and "Outlier" are classified as forgeries, empowering the detector to adapt effectively to unknown deepfakes.

We summarize FakeRadar's principal contributions as follows, with its conceptual pipeline illustrated in Figure 1:

- We propose FakeRadar, a novel framework for deepfake video detection, designed to address the critical challenge of cross-domain generalization in real-world scenarios.
- We introduce Forgery Outlier Probing for synthesizing outlier samples to cover a broader range of potential forger patterns, expanding the detector's understanding of forgery distribution.
- We design Outlier-Guided Tri-Training for model optimization, which enables end-to-end training for distinguishing real, fake, and outlier samples, improving detection of unknown forgeries.
- Numerous experiments have revealed that FakeRadar outperforms current leading algorithms across multiple deepfake benchmarks, particularly excelling in cross-domain evaluations where existing methods struggle.

2. Related work

Deepfake Detection. In response to the malicious use of deepfake techniques, numerous approaches have been explored to address this issue. Generally, deepfake detection is formulated as a binary classification problem (real or fake), leading researchers to develop a variety of end-to-end detectors [45, 54] that directly distinguish reals from manipulated content. Existing detection methods can be broadly categorized into two groups: image-level and video-level detection. Specifically, image-level detectors primarily focus on identifying spatial artifacts caused by forgeries. Examples include frequency anomalies [3, 16, 29, 39], identity-based discrepancies [13, 14, 24], blending artifacts [30, 44], and localized distortions [35, 49]. In contrast, video-level detectors employ a wider range of strategies. For example, some algorithms, such as [20, 32], concentrate on temporal consistency, typically reflected in correlations across consecutive frames. Other video detection studies [18, 38, 61] incorporate both temporal and spatiotemporal features to establish a more comprehensive detection framework. Another line of research [22, 56] examines audio-visual con-

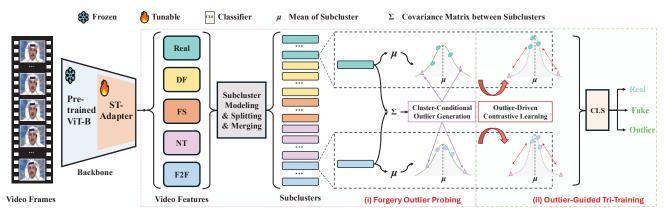


Figure 2. An overview of FakeRadar for deepfake video detection. FakeRadar consists of two key components: Forgery Outlier Probing and Outlier-Guided Tri-Training. (i) Forgery Outlier Probing performs dynamic subcluster modeling to model the feature-space distributions of real and fake samples using GMMs, followed by dynamic subcluster splitting and merging to produce finer-grained subcluster partitions. Afterwards, FakeRadar conducts cluster-conditioned outlier generation to simulate unseen forgeries for novel manipulation techniques. (ii) Outlier-Guided Tri-Training jointly optimizes the CLIP model's ViT-B/16 backbone, inserted with ST-Adapter [37] for parameter-efficient fine tuning, and a triplet-class classifier to distinguish between "Real", "Fake", and "Outlier" classes, with "Fake" and "Outlier" merged into a single new "Fake" class during evaluation. The training set includes real videos, as well as deepfake videos generated using four different manipulations, i.e. "DF" [41], "FS" [26], "NT" [47] and "F2F" [46].

sistency by analyzing the synchronization and correlation between video and audio signals to detect forgeries. In summary, these methods rely heavily on existing artifacts in training set, limiting their generalization to unseen manipulation types. In contrast, our approach explores featurespace anomalies, enhancing the generalization capacity of our proposed detector to diverse deepfake manipulations.

Towards Generalization in Deepfake Detection. Rapid advance in deepfakes have significantly improved the quality and authenticity of generated content, posing substantial challenges to the generalization capability of detection systems. Existing detectors [6, 43, 50] often suffer from severe performance degradation when encountering novel forgery techniques or previously unseen manipulations. To address this issue, researchers have proposed various strategies to enhance detector generalization. For instance, FTCN [61] improves generalization by capturing temporal coherence features in video sequences. Similarly, AltFreezing [51] employs an alternating weight-freezing training strategy to facilitate the joint detection of spatial and temporal forgery artifacts. Another algorithm, StyleFlow [5], leverages the temporal dynamics of style latent vectors to identify anomalous patterns in generated videos.

More recently, data synthesis has emerged as an effective strategy for improving detector generalization. For example, FWA [31] enhances training data by introducing facial warping artifacts, while Face X-ray [30] detects forgeries by identifying blending boundaries, independent of specific facial manipulation techniques. Building on this concept, SBI [44] synthesizes training samples by blending different views of the same pristine image to simulate generic forgery

artifacts. Additionally, a series of recent studies [1, 27, 34] have further advanced this direction by employing various data synthesis techniques. Unlike existing methods that primarily focus on image- or video-level data synthesis for deepfake detection, our approach proposes to explore unseen forgeries in feature space.

3. Methodology

3.1. Preliminaries

Overview. In this paper, we introduce *FakeRadar*, to overcomes the challenges of cross-domain generalization in deepfake video detection by exploring the feature space to identify emerging manipulation patterns. Specifically, FakeRadar is built upon two key components: (i) Forgery Outlier Probing, which models the feature-space distributions of subclusters for training samples of both real and fake categories, and then generates cluster-conditioned outlier samples to simulate unseen forgeries, and (ii) Outlier-Guided Tri-Training, which jointly optimizes an adapter-inserted backbone and a triplet-class classifier to distinguish among samples from the "Real", "Fake", and "Outlier" classes. An overview of FakeRadar can be found in Figure 2.

Model Adaptation with a Frozen Backbone. The FakeRadar framework employs a pre-trained Vision Transformer (ViT-B) from CLIP [40] as a frozen backbone to preserve its learned semantic features. This model serves as the video encoder for input clips. To adapt it for deepfake detection, we integrate lightweight ST-Adapter layers [37] after each Transformer block, enabling parameter-efficient finetuning. This significantly enhances the model's ability to

capture spatial manipulations and temporal inconsistencies, which are crucial for detecting deepfake videos. Given a feature vector x from the video encoder, each adapter transforms it as:

$$ST-Adapter(x) = x + ReLU(Conv3D(xW_{down}))W_{up}$$
. (1)

Here, $\mathrm{Conv3D}(\cdot)$ denotes a 3D convolution operation employed for capturing spatio-temporal features from videos. As well, W_{down} and W_{up} are weight matrices that perform dimensionality reduction and expansion, respectively, and $\mathrm{ReLU}(\cdot)$ is the activation function. By introducing only a small number of additional parameters, these adapters allow efficient fine-tuning while retaining the original backbone's representational strength. After feature extraction, each video clip is represented as a d-dimensional vector for subsequent analysis, with d=768.

Inference. During inference, the model classifies input samples as either real or deepfake, with deepfakes categorized as "Fake" or "Outlier". The "Fake" class includes fake samples with known manipulation patterns, while the "Outlier" class detects testing samples with novel forgeries that differ from established manipulations. By unifying these categories, the model enhances its robustness against evolving manipulation techniques, enabling accurate detection of both familiar and emerging deepfakes. This improves its generalization performance in cross-domain evaluations.

3.2. Forgery Outlier Probing

In this section, we introduce *Forgery Outlier Probing* (FOP) to generate outlier samples with unseen forgeries beyond the existing real and fake videos, simulating previously unobserved manipulation patterns. Specifically, FOP first presents dynamic subcluster modeling, which performs category-specific clustering on the training samples of real and fake categories. Then, it models the distributions of these categories using a strategy based on Gaussian Mixture Models (GMMs), and follows by applying dynamic subcluster splitting and merging, thereby achieving finer-grained subcluster partitioning. Finally, we generate cluster-conditioned outliers based on the resulting subclusters to stimulate unseen forgeries.

Dynamic Subcluster Modeling. To capture subtle intraclass variations within real and deepfake videos, we apply category-specific clustering on the FaceForensics++ (FF++) dataset. Specifically, we treat the *real* class and each of the four *fake* classes (i.e. four types of deepfake videos created by different manipulation techniques [26, 41, 46, 47]) as five separate categories. For each category c, let $\mathcal{X}^{(c)} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$ denote the set of feature embeddings (i.e. extracted from the video encoder). Inspired by [42], we then

define a main clustering network:

$$f_{\text{main}}^{(c)}(\mathbf{x}_i) \mapsto \mathbf{r}_i = (r_{i,1}, r_{i,2}, \dots, r_{i,K}),$$
 (2)

which initially assumes K clusters, empirically with K=5, providing a "soft" assignment of each embedding $\mathbf{x}_i \in \mathcal{X}^{(c)}$ to these K clusters. Here, $r_{i,k} \in [0,1]$ and $\sum_{k=1}^K r_{i,k} = 1$.

Each cluster can be modeled as a Gaussian distribution using a GMM, allowing us to represent the feature distribution of each category in a probabilistic manner. Specifically, for each cluster k in category c, we define a Gaussian component characterized by $\{\pi_k, \mu_k, \Sigma_k\}$, where $\pi_k \in (0,1)$ is the mixture weight for cluster k, with $\sum_{k=1}^K \pi_k = 1$, while μ_k and Σ_k represent its corresponding mean and covariance matrix. These parameters enable the clustering network to estimate the likelihood of each data sample belonging to different clusters, aligning our model with a GMM-based formulation.

After that, the main clustering network needs to learn these parameters by aligning its soft assignments with GMM responsibilities:

$$r_{i,k}^{E} = \frac{\pi_k \mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_{k'=1}^{K} \pi_{k'} \mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_{k'}, \boldsymbol{\Sigma}_{k'})},$$
(3)

where $\mathcal{N}(\cdot; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ is the multivariate Gaussian density function. The network is trained using the KL-divergence loss:

$$\mathcal{L}_{\text{main}} = \sum_{i=1}^{N} \text{KL}(\mathbf{r}_i \parallel \mathbf{r}_i^E), \tag{4}$$

which ensures that the learned cluster assignments match the inferred GMM responsibilities.

To further refine cluster structures, we introduce a subclustering network for each cluster k, which attempts to split the cluster into two subclusters if multi-modal structures are detected. A compactness loss is applied to encourage subclusters to be well-separated. This loss function penalizes subclusters that are too close to each other:

$$\mathcal{L}_{\text{sub}} = \sum_{k=1}^{K} \sum_{j=1}^{2} \sum_{\mathbf{x}_{i} \in X_{k}} \tilde{r}_{i,j} \| \mathbf{x}_{i} - \tilde{\boldsymbol{\mu}}_{k,j} \|_{2}^{2},$$
 (5)

where $\tilde{r}_{i,j}$ represents the probability of \mathbf{x}_i belonging to subcluster j, and $\tilde{\mu}_{k,j}$ is its mean.

Since the predefined cluster number K may not be optimal, we propose to dynamically adjust it through split/merge proposals, inspired by [42]. Suppose a cluster k contains N_k points, X_k . We *split* if subclustering indicates that X_k is better viewed as two subgroups, $\{X_{k,1}, X_{k,2}\}$. The Hastings ratio is given by

$$H_s = \frac{\alpha \Gamma(N_{k,1}) f_x(X_{k,1}) \Gamma(N_{k,2}) f_x(X_{k,2})}{\Gamma(N_k) f_x(X_k)}, \quad (6)$$

where $f_x(\cdot)$ is the marginal likelihood under a weak NIW prior, α is the Dirichlet-process concentration parameter, and $\Gamma(\cdot)$ is the Gamma function.

If H_s exceeds a threshold, cluster k is split into two new clusters. The parameters for these new clusters are initialized from $\tilde{\mu}_{k,1}$, $\tilde{\Sigma}_{k,1}$ and $\tilde{\mu}_{k,2}$, $\tilde{\Sigma}_{k,2}$, respectively. Merging is similarly proposed if two clusters significantly overlap, with acceptance ratio $H_m = 1/H_s$. Repeating these steps yields an effective number of clusters $K^{(c)}$, plus subclusters for each category c.

After these dynamic adjustments, we obtain the final subcluster set $\hat{\mathcal{C}} = \{\hat{C}_1, \hat{C}_2, ..., \hat{C}_{K^{(c)}}\}$, where $K^{(c)}$ is the refined number of clusters. Each adjusted subcluster \hat{C}_k is characterized by its updated mean and covariance:

$$\hat{\boldsymbol{\mu}}_k = \frac{1}{|\hat{C}_k|} \sum_{\mathbf{x}_i \in \hat{C}_k} \mathbf{x}_i,\tag{7}$$

$$\hat{\Sigma} = \frac{1}{|\hat{C}_k|} \sum_{\mathbf{x}_i \in \hat{C}_k} (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_k) (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_k)^{\top}.$$
 (8)

We use online estimation for efficient training by maintaining a cluster-conditional queue with $|\hat{C}_k|$ sample instances from each subcluster. In each iteration, we enqueue the embeddings of objects to their corresponding cluster-conditional queues and dequeue the same number of object embeddings.

Cluster-Conditional Outlier Generation. After deriving finer-grained cluster partitions for each category, we generate feature-space outliers to probe potential unknown forgeries. Motivated by [15], in order to ensure that the sampled outliers lie near the cluster boundary, we maintain a sufficiently small value ϵ for outlier sampling from the ϵ -likelihood region of the estimated cluster-conditional distribution. Therefore, for a subcluster from $\hat{\mathcal{C}}$ indexed by k, the outlier samples it generates can be formulated as follows:

$$\mathcal{V}_{k} = \left\{ \mathbf{v}_{k} \mid \frac{\exp\left(-\frac{1}{2}(\mathbf{v}_{k} - \hat{\boldsymbol{\mu}}_{k})^{\top} \hat{\boldsymbol{\Sigma}}^{-1}(\mathbf{v}_{k} - \hat{\boldsymbol{\mu}}_{k})\right)}{(2\pi)^{d/2} |\hat{\boldsymbol{\Sigma}}|^{1/2}} < \epsilon \right\},$$
(9)

where $v_k \sim \mathcal{N}(\hat{\mu}_k, \hat{\Sigma}_k)$ represents the virtual outliers sampled from the estimated Gaussian distribution of subcluster k, with d denoting its embedding dimensionality. It is worth noting that during model optimization, all outliers generated here, i.e. $\overline{\mathcal{V}} = \mathcal{V}_1 \cup \cdots \cup \mathcal{V}_k \cup \cdots \mathcal{V}_{|\hat{\mathcal{C}}|}$, are classified into the "Outlier" class, separate from the real and fake classes.

3.3. Outlier-Guided Tri-Training

To enable better performance of FakeRadar for tasks of deepfake video detection, we design a strategy called *Outlier-Guided Tri-Training* for model optimization jointly

combining proposed outlier-driven contrastive loss and outlier-conditioned cross-entropy loss. With this approach, a triplet-class classifier (integrated with the pre-trained backbone and adapters) is trained to explicitly distinguish among three categories: "Real", "Fake", and "Outlier". This design not only improves the model's ability to detect known forgeries but also sensitizes it to emerging manipulation patterns.

Outlier-Driven Contrastive Loss. In this work, we propose an outlier-driven contrastive loss to enhance the model's ability to delineate class boundaries within the feature space, thereby improving feature discriminability. This loss function is designed to maximize the similarity between samples and their respective subcluster centers while minimizing their similarity to other subcluster centers. Additionally, it treats generated outliers as negative examples, enforcing low similarity between samples and outliers to strengthen boundary discrimination.

To this end, the loss is formulated using an InfoNCE-like contrastive loss given an input sample i:

$$\mathcal{L}_{\text{con}} = -\log \frac{\exp\left(\frac{s(h_i, \mu_{k^+})}{\tau}\right)}{\sum_{k^-} \exp\left(\frac{s(h_i, \mu_{k^-})}{\tau}\right) + \sum_{v_j \in \overline{\mathcal{V}}} \exp\left(\frac{s(h_i, v_j)}{\tau}\right)},$$
(10)

where h_i denotes the feature representation of sample i, while $\mathcal V$ represents the outlier cache, providing additional hard negative samples to improve the model's ability to distinguish boundary samples. Additionally, k^- represents all subclusters except the one to which sample i belongs, with its corresponding u_{k^-} denoting the centers of these subclusters, used as negative samples. The function $s(\cdot,\cdot)$ denotes cosine similarity, which quantifies the similarity between samples, and τ is a temperature hyperparameter that controls the smoothness of similarity scaling.

Outlier-Conditioned Cross-Entropy Loss. While the proposed outlier-driven contrastive loss enforces stronger feature-space separation, we still require a concrete classification objective to handle a three-category task: real, fake, and outlier. In particular, the model must not only determine whether a video is fake or real but also assign a separate label for those samples situated near decision boundaries (i.e., outliers). This triplet-class cross-entropy loss is crucial, as it enables the model to explicitly differentiate uncertain boundary samples and label them as outliers, thereby avoiding confusion.

Concretely, the triplet-class cross-entropy loss can be defined as follows:

$$\mathcal{L}_{\text{cls}} = -\sum_{c \in \{\text{Real, Fake, Outlier}\}} y_c \log p_c, \tag{11}$$

Method	Backbone	Input Type	Testing Set AUC (%)					
Troutou	Buckeone		FF++	CDFv2	DFDCP	DFDC	DFD	
Xception [6]	Xception	Frame	99.3	73.7	-	70.9	-	
Face X-Ray [30]	HRNet	Frame	97.8	79.5	-	65.5	95.4	
PCL+I2G [60]	Resnet	Frame	99.8	90.0	74.4	67.5	-	
SLADD [4]	Xception	Frame	98.4	79.7	76.0	-	-	
UIA-ViT [63]	Vision Transformer	Frame	99.3	82.4	75.8	-	94.7	
UCF [53]	Xception	Frame	-	82.4	-	80.5	94.5	
SeeABLE [27]	EfficientNet	Frame	-	87.3	86.3	75.9	-	
LSDA [54]	EfficientNet	Frame	-	91.1	81.2	77.0	95.6	
FTCN [61]	3D ResNet	Video	99.8	86.9	-	74.0	-	
RealForensics [21]	3D ResNet	Video	99.8	86.9	-	75.9	-	
LTTD [19]	Vision Transformer	Video	99.5	89.3	-	80.4	-	
AltFreezing [51]	3D ResNet	Video	99.7	89.5	-	-	93.7	
TALL [52]	Swin Transformer	Video	99.9	90.8	-	76.8	-	
StyleFlow [5]	3D ResNet	Video	99.1	89.0	-	-	96.1	
NACO [59]	Vision Transformer	Video	99.8	89.5	-	76.7	-	
FakeRadar (Ours)	Vision Transformer	Video	99.1	91.7	88.5	84.1	96.2	

Table 1. Generalization comparisons of FakeRadar and state-of-the-art algorithms on **cross-dataset evaluation**, using video-level AUC (%) as the metric.

where y_c is the ground-truth label for samples belonging to class c, and p_c is the model-predicted probability of samples from that class, with c= "Outlier" denoting the outlier label. This loss compels the model to accurately partition each category in the feature space, especially for samples near class boundaries.

Total Loss. Finally, we combine both losses into the total objective for model optimization:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{con}} + \lambda \, \mathcal{L}_{\text{cls}},$$
 (12)

where $\lambda=0.5$ is a hyperparameter balancing the relative importance of classification and contrastive goals. In practice, each mini-batch comprises 16 training samples (FF++), and 16 outlier samples during model training.

4. Experiments

4.1. Experiment Setups

Datasets. Following established experimental practices [5, 51, 52, 54], this study primarily utilizes the high-quality (HQ) version of the FaceForensics++ dataset [43], denoted as FF++(HQ), for FakeRadar's evaluation. This dataset consists of 1000 pristine videos and 4000 deepfakes videos generated using four different manipulation techniques: DFD [41], F2F [46], FS [26], and NT [47]. The dataset is split into training, validation, and testing sets according to the official protocol. This study primarily focuses on evaluating the cross-domain performance of our method. To this end, we incorporate four additional datasets for evaluation, including CDFv2 [33], DFDCP [11], DFDC [12], and DFD [41].

Implementation Details. To ensure fair comparisons, we maintain most implementation details consistent with exist-

Training Set	Method	DF	F2F	FS	NT	avg.
	DCL [45]	99.98	77.13	61.01	75.01	78.28
DF	UIA-ViT [63]	99.29	74.44	53.89	70.92	74.64
	TALL [52]	99.35	69.33	54.38	74.74	74.45
	AltFreezing [51]	98.80	69.33	72.07	77.41	79.40
	FakeRadar (Ours)	99.95	86.47	88.12	66.36	82.23
F2F	DCL [45]	91.91	99.21	59.58	66.67	79.34
	UIA-ViT [63]	83.39	98.32	68.37	67.17	79.31
	TALL [52]	85.52	98.67	62.15	70.24	79.15
	AltFreezing [51]	81.12	98.80	62.97	70.26	78.29
	FakeRadar (Ours)	96.94	99.09	82.33	68.46	86.75
FS	DCL [45]	74.80	69.75	99.90	52.60	74.26
	UIA-ViT [63]	81.02	66.30	99.04	49.26	73.91
	TALL [52]	77.83	67.30	98.77	52.57	74.12
	AltFreezing [51]	80.67	65.94	99.78	57.23	75.90
	FakeRadar (Ours)	96.39	82.06	99.83	52.64	82.73
NT	DCL [45]	91.23	52.13	79.31	98.97	80.41
	UIA-ViT [63]	79.37	67.89	45.94	94.59	71.95
	TALL [52]	84.02	72.85	51.66	95.16	75.92
	AltFreezing [51]	88.81	76.78	57.92	95.71	79.80
	FakeRadar (Ours)	96.71	83.34	67.50	98.41	86.49

Table 2. Generalization comparisons of FakeRadar and state-of-the-art algorithms on **cross-manipulation evaluation**, using video-level AUC (%) as the metric.

ing methods [5, 51, 59]. However, our FakeRadar framework employs ViT-B/16 from CLIP [40] as the backbone, into which we insert an ST-Adapter [37] with a 384-dimensional bottleneck for parameter-efficient fine-tuning. Additionally, the classification head of our triplet-class classifier consists of a single linear layer with softmax activation. During model training, we sample four temporal clips from each video, each containing 12 consecutive frames. We use Adam as the optimizer and apply a cosine learning rate scheduler with an initial learning rate of 1×10^{-4} . The model is trained for 60 epochs with a batch size of 16 and a weight decay of 5×10^{-4} . Prior to model optimization, we perform data preprocessing by detecting faces with RetinaFace [10] and then standardize all face images

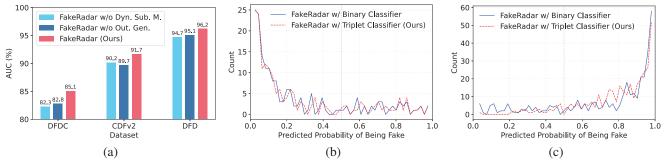


Figure 3. Evaluation of (a) Forgery Outlier Probing and (b)/(c) Outlier-Guided Tri-Training by FakeRadar. (a) "FakeRadar w/o Dyn. Sub. M." and "FakeRadar w/o Out. Gen." indicate that dynamic subcluster modeling and cluster-conditional outlier generation have been removed. The model is trained on FF++(HQ) and evaluated on DFDC [12], CDFv2 [33], and DFD [41]. (b) and (c) show, respectively, the count of real and fake test samples classified as "fake" across predicted probability intervals. The blue solid line denotes the binary classifier's distribution (distinguishing only real vs. fake), while the red dashed line denotes that of our proposed triplet classifier, which treats both fake and outlier samples as the new "fake". In this evaluation, the model is trained on FF++(HQ) and tested on CDFv2 [33].

to 224×224 pixels.

Evaluation Metrics. In this work, consistent with previous studies [5, 51, 52] employing video-level inputs, we use the area under the ROC curve (AUC) as the performance metric during evaluations. For methods that use images as model inputs, e.g. [27, 54], their video-level results are derived by averaging the predictions across all video frames.

4.2. Comparisons with the State-of-the-Arts

In this study, we focus on evaluating the capabilities of cross-domain generalization for our detector to validate its effectiveness in detecting deepfake videos manipulated through various techniques. To achieve this, we employ two strategies for evaluation: cross-dataset evaluation and cross-manipulation evaluation, both of which demonstrate that our results predominantly surpass existing state-of-theart (SoTA) algorithms for deepfake video detection, as illustrated by Table 1 and Table 2.

Cross-Dataset Evaluation. In this situation, we train our FakeRadar on the FF++(HQ) and test its performance on multiple benchmark datasets, including CDFv2 [33], DFD [41], DFDC [12], and DFDCP [11]. As reported in Table 1, FakeRadar demonstrates its superior generalization capability in cross-dataset evaluation, achieving the best performance across all cross-dataset cases while maintaining comparable performance on the FF++ test dataset against SoTA algorithms. Specifically, our algorithm significantly improves over other methods with images and videos as inputs, particularly outperforming the current best methods, UCF [53] and LTTD [19], by 3.6% and 3.7% in the AUC metric, respectively, on the dataset DFDC [12]. These results robustly showcase the proposed method's enhanced ability of generalization to handle unseen manipulation techniques in cross-dataset evaluation scenarios.

M-(#)	FOP	ODCL	TT OCCE	CDF	DFDC	DFDCP	DFD	avg.
1	/	✓	✓	91.6	84.1	88.5	96.2	90.1
2	✓	✓		89.9	81.2	85.6	94.9	87.9
3	✓		✓	88.6	80.9	87.7	94.7	88.0
4		✓		88.8	80.2	86.7	94.5	87.6
5			✓	88.4	78.4	85.1	94.3	86.7
6				88.2	78.3	84.8	94.2	86.4

Table 3. Ablation study results of FakeRadar by removing its one or more proposed components, using video-level AUC (%) as the metric. For clarity, "FOP", "OGTT", "ODCL", and "OCCE" refer to Forgery Outlier Probing, Outlier-Guided Tri-Training, Outlier-Driven Contrastive Loss, and Outlier-Conditioned Cross-Entropy Loss, respectively.

Cross-Manipulation Evaluation. To further assess our model's generalization across different manipulation techniques, we conduct cross-manipulation evaluation on the dataset FF++(HQ). In this case, we train the model using only real samples from FF++(HQ) and one manipulation type of deepfake samples, then test its performance on the remaining three manipulation types within FF++(HQ). As shown in Table 2, the proposed FakeRadar achieves the best average performance in all testing cases, particularly excelling in scenarios trained with F2F as the manipulation type, surpassing the second-best performing method, DCL [45], by 7.41% in AUC. In evaluations focused on a single manipulation type, our method also achieves the best performance in most cases and the second-best in a few. This indicates that our approach, in cross-manipulation evaluation, demonstrates a performance advantage in crossdomain generalization compared to current SoTA algorithms of detecting deepfake videos.

4.3. Ablation Analysis

We conduct an ablation study to assess the contribution of each component in FakeRadar, with results presented in Table 3, Figure 3, and Figure 4. Experiments are performed in a cross-dataset setting, where the model is trained on FF++(HQ) and tested on CDF, DFDC, DFDCP, and DFD, with average performance reported. As shown in Table 3 (Rows M-(1) to M-(6)), removing individual or multiple components leads to a progressive decline in AUC, from 90.1% to 86.4%, demonstrating the significance of each component. Below, we analyze the impact of FakeRadar's key components in detail.

Effectiveness of Forgery Outlier Probing. To evaluate the impact of FOP, we remove or modify its submodules: dynamic subcluster modeling and cluster-conditional outlier generation. In "FakeRadar w/o Dyn. Sub. M.", we disable dynamic subcluster modeling, generating outliers only from clusters of samples of real and fake categories in training set. In "FakeRadar w/o Out. Gen.", we replace outlier samples with Gaussian noise. As shown in Figure 3(a), removing dynamic subcluster modeling reduces AUC on CDFv2, DFDC, and DFD to 90.2%, 82.3%, and 94.7%, respectively. While removing outlier generation has a smaller impact, performance remains below the complete FakeRadar framework. On CDFv2, FakeRadar (91.7%) outperforms "FakeRadar w/o Out. Gen." by 2.0%, demonstrating that FOP enhances generalization by generating effective outlier samples that improve the detection of novel deepfake manipulations.

Effectiveness of Outlier-Guided Tri-Training. Comparing Table 3 (Rows M-(1) to M-(3)), removing Outlier-Driven Contrastive Loss or Outlier-Conditioned Cross-Entropy Loss leads to an AUC drop from 90.1% to 87.9% and 88.9%, respectively, highlighting their importance. Since both losses refine the use of outlier samples, removing them would render FOP ineffective.

The core of our proposed optimization strategy is the integration of a triplet classifier to handle the outlier samples generated by FOP. To assess its impact, we replace the triplet classifier, which performs three-class classification (real, fake, and outlier), with a binary classifier that only distinguishes between real and fake. We then analyze the prediction probability distributions for cross-dataset evaluation, as shown in Figure 3 (b) and (c). As illustrated, for real samples, both the binary and triplet classifiers exhibit similar prediction patterns with high confidence in the lowprobability region. However, for fake samples, the triplet classifier demonstrates a clear advantage. Compared to the binary classifier, it assigns higher forgery probabilities to a greater number of fake samples, indicating that the proposed outlier-guided tri-training increases the model's confidence in detecting unseen fake samples. This result validates the effectiveness of treating outlier samples as fake samples during inference and highlights its contribution to improving the model's generalization capability.

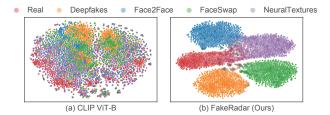


Figure 4. Feature visualization using t-SNE. We visualize the features of (a) the pre-trained CLIP ViT-B model and our (b) FakeR-adar here. We train both models using the training set FF++(HQ), while testing on the FF++ test set.

Feature Visualization. We use t-SNE [48] to visualize the features of our FakeRadar and the pre-trained CLIP ViT-B model. As shown in Figure 4, the feature distribution of FakeRadar exhibits a more compact clustering structure and clearer boundaries between real and fake videos, compared to the pre-trained CLIP ViT-B model. Furthermore, FakeRadar effectively distinguishes real samples from the four types of forgery samples. This suggests that FakeRadar not only enables the model to learn discriminative features between categories but also establishes more robust decision boundaries with the aid of outlier samples. These feature distribution patterns underscore the advantages of FakeRadar in improving the model's generalization ability. The boundaries between categories are sharper, and sufficient tolerance space is provided, which is essential for addressing unseen forgery types in real-world scenarios.

5. Conclusions

In this work, we introduced FakeRadar, a novel deepfake video detection framework designed to address the limitations of existing methods in cross-domain generalization. Unlike traditional detectors that passively learn specific forgery patterns, FakeRadar proactively probes unknown forgeries by incorporating Forgery Outlier Probing to simulate diverse unseen manipulation patterns. This approach enables the model to explore a broader range of potential forgeries in feature space, significantly enhancing its adaptability to novel manipulations. To further improve detection robustness, we proposed Outlier-Guided Tri-Training, allowing the model to distinguish between real, fake, and outlier samples. This training paradigm, introduced outlier-driven contrastive loss and outlier-conditioned cross-entropy loss, sharpens the decision boundaries and enhances the detector's capacity to generalize across different manipulation types. Extensive experiments on benchmark datasets confirm that FakeRadar consistently outperforms state-of-the-art detection algorithms, particularly in cross-domain evaluations where conventional detectors struggle.

Acknowledgments

This work was supported in part by the National Key R&D Program of China (2024YFB3908503), in part by the Guangxi Science and Technology Project (AB25069496), and in part by the National Natural Science Foundation of China (62172120 and 62322608).

References

- [1] Weiming Bai, Yufan Liu, Zhipeng Zhang, Bing Li, and Weiming Hu. Aunet: Learning relations between action units for face forgery detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 24709–24719, 2023. 3
- [2] Junyi Cao, Chao Ma, Taiping Yao, Shen Chen, Shouhong Ding, and Xiaokang Yang. End-to-end reconstructionclassification learning for face forgery detection. In 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 4103–4112, 2022. 2
- [3] Keshigeyan Chandrasegaran, Ngoc-Trung Tran, and Ngai-Man Cheung. A closer look at fourier spectrum discrepancies for cnn-generated images detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7200–7209, 2021. 2
- [4] 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. 6
- [5] Jongwook Choi, Taehoon Kim, Yonghyun Jeong, Seungryul Baek, and Jongwon Choi. Exploiting style latent flows for generalizing deepfake video detection. In *Proceedings of* the IEEE/CVF conference on computer vision and pattern recognition, pages 1133–1143, 2024. 2, 3, 6, 7
- [6] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE con*ference on computer vision and pattern recognition, pages 1251–1258, 2017. 3, 6
- [7] Danielle Keats Citron and Robert M. Chesney. Deepfakes and the new disinformation war. *Foreign Affairs*, 2018. 1
- [8] Di Cooke, Abigail Edwards, Sophia Barkoff, and Kathryn Kelly. As good as a coin toss: Human detection of aigenerated images, videos, audio, and audiovisual stimuli. *ArXiv*, abs/2403.16760, 2024. 1
- [9] Yunshu Dai, Jianwei Fei, and Fangjun Huang. Idguard: Robust, general, identity-centric poi proactive defense against face editing abuse. In 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 11934– 11943, 2024. 1
- [10] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 5202–5211, 2020. 6
- [11] B Dolhansky. The dee pfake detection challenge (dfdc) pre view dataset. arXiv preprint arXiv:1910.08854, 2019. 2, 6, 7

- [12] Brian Dolhansky, Joanna Bitton, Ben Pflaum, Jikuo Lu, Russ Howes, Menglin Wang, and Cristian Canton Ferrer. The deepfake detection challenge (dfdc) dataset. arXiv preprint arXiv:2006.07397, 2020. 2, 6, 7
- [13] Chengdong Dong, Ajay Kumar, and Eryun Liu. Think twice before detecting gan-generated fake images from their spectral domain imprints. In 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 7855–7864, 2022. 2
- [14] Shichao Dong, Jin Wang, Renhe Ji, Jiajun Liang, Haoqiang Fan, and Zheng Ge. Implicit identity leakage: The stumbling block to improving deepfake detection generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3994–4004, 2023. 2
- [15] Xuefeng Du, Zhaoning Wang, Mu Cai, and Sharon Li. Towards unknown-aware learning with virtual outlier synthesis. In *International Conference on Learning Representa*tions, 2022. 5
- [16] Joel Frank, Thorsten Eisenhofer, Lea Schönherr, Asja Fischer, Dorothea Kolossa, and Thorsten Holz. Leveraging frequency analysis for deep fake image recognition. In *International conference on machine learning*, pages 3247–3258. PMLR, 2020. 2
- [17] Qiqi Gu, Shen Chen, Taiping Yao, Yang Chen, Shouhong Ding, and Ran Yi. Exploiting fine-grained face forgery clues via progressive enhancement learning. In AAAI Conference on Artificial Intelligence, 2021. 1
- [18] Zhihao Gu, Yang Chen, Taiping Yao, Shouhong Ding, Jilin Li, Feiyue Huang, and Lizhuang Ma. Spatiotemporal inconsistency learning for deepfake video detection. In *Proceedings of the 29th ACM international conference on multimedia*, pages 3473–3481, 2021. 2
- [19] Jiazhi Guan, Hang Zhou, Zhibin Hong, Errui Ding, Jingdong Wang, Chengbin Quan, and Youjian Zhao. Delving into sequential patches for deepfake detection. Advances in Neural Information Processing Systems, 35:4517–4530, 2022. 6, 7
- [20] Alexandros Haliassos, Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Lips don't lie: A generalisable and robust approach to face forgery detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5039–5049, 2021. 2
- [21] Alexandros Haliassos, Rodrigo Mira, Stavros Petridis, and Maja Pantic. Leveraging real talking faces via selfsupervision for robust forgery detection. In *Proceedings of* the IEEE/CVF conference on computer vision and pattern recognition, pages 14950–14962, 2022. 6
- [22] Alexandros Haliassos, Rodrigo Mira, Stavros Petridis, and Maja Pantic. Leveraging real talking faces via selfsupervision for robust forgery detection. In *Proceedings of* the IEEE/CVF conference on computer vision and pattern recognition, pages 14950–14962, 2022. 2
- [23] Ammarah Hashmi, Sahibzada Adil Shahzad, Chia-Wen Lin, Yu Tsao, and Hsin-Min Wang. Unmasking illusions: Understanding human perception of audiovisual deepfakes. ArXiv, abs/2405.04097, 2024. 1
- [24] Baojin Huang, Zhongyuan Wang, Jifan Yang, Jiaxin Ai, Qin Zou, Qian Wang, and Dengpan Ye. Implicit identity

- driven deepfake face swapping detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4490–4499, 2023. 2
- [25] 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. 1
- [26] Marek Kowalski. Faceswap. https://github.com/ MarekKowalski/FaceSwap, 2018. GitHub repository, accessed on March 3, 2025. 3, 4, 6
- [27] Nicolas Larue, Ngoc-Son Vu, Vitomir Struc, Peter Peer, and Vassilis Christophides. Seeable: Soft discrepancies and bounded contrastive learning for exposing deepfakes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 21011–21021, 2023. 1, 3, 6, 7
- [28] Hong Li, Yutang Feng, Song Xue, Xuhui Liu, Bohan Zeng, Shanglin Li, Boyu Liu, Jianzhuang Liu, Shumin Han, and Baochang Zhang. Uv-idm: Identity-conditioned latent diffusion model for face uv-texture generation. In 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 10585–10595, 2024. 1
- [29] Hanzhe Li, Jiaran Zhou, Yuezun Li, Baoyuan Wu, Bin Li, and Junyu Dong. Freqblender: Enhancing deepfake detection by blending frequency knowledge. *arXiv preprint arXiv:2404.13872*, 2024. 2
- [30] 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. 1, 2, 3, 6
- [31] Yuezun Li and Siwei Lyu. Exposing deepfake videos by detecting face warping artifacts. arXiv preprint arXiv:1811.00656, 2018. 3
- [32] Yuezun Li, Ming-Ching Chang, and Siwei Lyu. In ictu oculi: Exposing ai created fake videos by detecting eye blinking. In 2018 IEEE International Workshop on Information Forensics and Security (WIFS), pages 1–7, 2018. 1, 2
- [33] Yuezun Li, Xin Yang, Pu Sun, Honggang Qi, and Siwei Lyu. Celeb-df: A large-scale challenging dataset for deepfake forensics. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 3204–3213, 2020. 6, 7
- [34] Yuzhen Lin, Wentang Song, Bin Li, Yuezun Li, Jiangqun Ni, Han Chen, and Qiushi Li. Fake it till you make it: Curricular dynamic forgery augmentations towards general deepfake detection. In *European Conference on Computer Vision*, pages 104–122. Springer, 2024. 3
- [35] Dat NGUYEN, Nesryne MEJRI, Inder Pal SINGH, Polina KULESHOVA, Marcella ASTRID, Anis KACEM, Enjie GHORBEL, and Djamila AOUADA. Laa-net: Localized artifact attention network for quality-agnostic and generalizable deepfake detection. In 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 17395–17405, 2024. 1, 2
- [36] Utkarsh Ojha, Yuheng Li, and Yong Jae Lee. Towards universal fake image detectors that generalize across generative models. In *Proceedings of the IEEE/CVF Conference*

- on Computer Vision and Pattern Recognition, pages 24480–24489, 2023. 2
- [37] Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hong-sheng Li. St-adapter: Parameter-efficient image-to-video transfer learning. Advances in Neural Information Processing Systems, 35:26462–26477, 2022. 3, 6
- [38] Guilin Pang, Baopeng Zhang, Zhu Teng, Zige Qi, and Jianping Fan. Mre-net: Multi-rate excitation network for deepfake video detection. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(8):3663–3676, 2023. 2
- [39] Yuyang Qian, Guojun Yin, Lu Sheng, Zixuan Chen, and Jing Shao. Thinking in frequency: Face forgery detection by mining frequency-aware clues. In *European conference on com*puter vision, pages 86–103. Springer, 2020. 2
- [40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 2021. 2, 3, 6
- [41] Md Shohel Rana, Mohammad Nur Nobi, Beddhu Murali, and Andrew H Sung. Deepfake detection: A systematic literature review. *IEEE access*, 10:25494–25513, 2022. 3, 4, 6, 7
- [42] Meitar Ronen, Shahaf E. Finder, and Oren Freifeld. Deepdpm: Deep clustering with an unknown number of clusters. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 9861–9870, 2022. 4
- [43] Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Niessner. Faceforensics++: Learning to detect manipulated facial images. In 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pages 1–11, 2019. 3, 6
- [44] 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. 1, 2, 3
- [45] Ke Sun, Taiping Yao, Shen Chen, Shouhong Ding, Jilin Li, and Rongrong Ji. Dual contrastive learning for general face forgery detection. In *Proceedings of the AAAI conference on artificial intelligence*, pages 2316–2324, 2022. 2, 6, 7
- [46] Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Face2face: Real-time face capture and reenactment of rgb videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2387–2395, 2016. 1, 3, 4, 6
- [47] 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. 1, 3, 4, 6
- [48] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9 (86):2579–2605, 2008. 8
- [49] 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 detec-

- tion. In Proceedings of the 2022 international conference on multimedia retrieval, pages 615–623, 2022. 2
- [50] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. Cnn-generated images are surprisingly easy to spot... for now. In *Proceedings of* the IEEE/CVF conference on computer vision and pattern recognition, pages 8695–8704, 2020. 3
- [51] Zhendong Wang, Jianmin Bao, Wengang Zhou, Weilun Wang, and Houqiang Li. Altfreezing for more general video face forgery detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4129–4138, 2023. 2, 3, 6, 7
- [52] Yuting Xu, Jian Liang, Gengyun Jia, Ziming Yang, Yanhao Zhang, and Ran He. Tall: Thumbnail layout for deepfake video detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 22658–22668, 2023. 6, 7
- [53] Zhiyuan Yan, Yong Zhang, Yanbo Fan, and Baoyuan Wu. Ucf: Uncovering common features for generalizable deep-fake detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 22412–22423, 2023. 6, 7
- [54] Zhiyuan Yan, Yuhao Luo, Siwei Lyu, Qingshan Liu, and Baoyuan Wu. Transcending forgery specificity with latent space augmentation for generalizable deepfake detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8984–8994, 2024. 1, 2, 6, 7
- [55] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. ACM Computing Surveys, 56(4): 1–39, 2023. 1
- [56] Wenyuan Yang, Xiaoyu Zhou, Zhikai Chen, Bofei Guo, Zhongjie Ba, Zhihua Xia, Xiaochun Cao, and Kui Ren. Avoid-df: Audio-visual joint learning for detecting deepfake. *IEEE Transactions on Information Forensics and Security*, 18:2015–2029, 2023. 2
- [57] Shenghai Yuan, Jinfa Huang, Xianyi He, Yunyuan Ge, Yujun Shi, Liuhan Chen, Jiebo Luo, and Li Yuan. Identitypreserving text-to-video generation by frequency decomposition, 2024. 1
- [58] Yu Zeng, Zhe L. Lin, and Vishal M. Patel. Sketchedit: Mask-free local image manipulation with partial sketches. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 5941–5951, 2021. 1
- [59] Daichi Zhang, Zihao Xiao, Shikun Li, Fanzhao Lin, Jianmin Li, and Shiming Ge. Learning natural consistency representation for face forgery video detection. In European Conference on Computer Vision, pages 407–424. Springer, 2024.
- [60] Tianchen Zhao, Xiang Xu, Mingze Xu, Hui Ding, Yuanjun Xiong, and Wei Xia. Learning self-consistency for deepfake detection. In 2021 IEEE/CVF International Conference on Computer Vision (ICCV), pages 15003–15013, 2021. 6
- [61] Yinglin Zheng, Jianmin Bao, Dong Chen, Ming Zeng, and Fang Wen. Exploring temporal coherence for more general video face forgery detection. In *Proceedings of the*

- IEEE/CVF international conference on computer vision, pages 15044–15054, 2021. 2, 3, 6
- [62] Yang Zhou, Zichong Chen, and Hui Huang. Deformable oneshot face stylization via dino semantic guidance. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 7787–7796, 2024. 1
- [63] Wanyi Zhuang, Qi Chu, Zhentao Tan, Qiankun Liu, Haojie Yuan, Changtao Miao, Zixiang Luo, and Nenghai Yu. Uiavit: Unsupervised inconsistency-aware method based on vision transformer for face forgery detection. In *European conference on computer vision*, pages 391–407. Springer, 2022.