

Fake It till You Make It: Curricular Dynamic Forgery Augmentations towards General Deepfake Detection

Yuzhen Lin ^{*1}, Wentang Song ^{*1}, Bin Li ^{✉1}, Yuezun Li², Jiangqun Ni³,
Han Chen¹, and Qiushi Li¹

¹ Guangdong Provincial Key Laboratory of Intelligent Information Processing,
Shenzhen Key Laboratory of Media Security, SZU-AFS Joint Innovation Center for
AI Technology, Shenzhen University, Shenzhen, China
{linyuzhen2020, 2018132120, 2016130205, 1800271017}@email.szu.edu.cn;

✉ libin@szu.edu.cn;

² College of Computer Science and Technology, Ocean University of China, Qingdao,
China; liyuezun@ouc.edu.cn

³ School of Cyber Science and Technology, Sun Yat-Sen University, and Department
of New Networks, Peng Cheng Laboratory, Shenzhen, China
issjqni@mail.sysu.edu.cn

Abstract. Previous studies in deepfake detection have shown promising results when testing face forgeries from the same dataset as the training. However, the problem remains challenging when one tries to generalize the detector to forgeries from unseen datasets and created by unseen methods. In this work, we present a novel general deepfake detection method, called **Curricular Dynamic Forgery Augmentation (CDFA)**, which jointly trains a deepfake detector with a forgery augmentation policy network. Unlike the previous works, we propose to progressively apply forgery augmentations following a monotonic curriculum during the training. We further propose a dynamic forgery searching strategy to select one suitable forgery augmentation operation for each image varying between training stages, producing a forgery augmentation policy optimized for better generalization. In addition, we propose a novel forgery augmentation named self-shifted blending image to simply imitate the temporal inconsistency of deepfake generation. Comprehensive experiments show that CDFA can significantly improve both cross-datasets and cross-manipulations performances of various naive deepfake detectors in a plug-and-play way, and make them attain superior performances over the existing methods in several benchmark datasets.

Keywords: Deepfake Detection · Curriculum Learning · Forgery Augmentation

* Equal contributions. ✉ Corresponding author.

1 Introduction

Deepfake techniques [18, 26, 39, 40, 55, 57] refer to a series of deep learning-based facial forgery techniques that can swap or reenact the face of one person in a video to another. It poses a significant threat given their potential by spreading false information and even political manipulation. To reduce these risks, detecting deepfakes has become a crucial research topic in recent years.

Early works [1, 43] treat deepfake detection as a binary classification problem and directly use deep neural networks [11, 50] to distinguish fake faces (named naive deepfake detectors [60]). In order to improve the detection performance, some works [6, 32, 35, 42, 62] introduce auxiliary modalities (e.g., frequency) or supervision (e.g., forgery masks) information for learning subtle forgery artifacts. These methods achieve promising performance in a closed-domain scenario, where the training and testing data are sampled from the same distribution. However, in practice the testing forgeries are usually from unseen datasets and synthesized by unseen methods. Discrepancies between training and testing data lead to inferior performance of detectors, which poses challenges to deepfake detectors for practical usage.

Recall that a forgery can be easily synthesized by blending two different images. Motivated by this, a powerful solution to improve the generalization capabilities of deepfake detectors is introducing the forgery augmentation technology [27, 45] that blends two real faces from training data to get new face forgeries. The augmented sample (labeled as fake) is so-called pseudo fake (*p-fake*) sample [20] to distinguish them from the original fake (*o-fake*) sample of the training data. Forgery augmentation strategies are also at the core of many state-of-the-art (SOTA) detection models [2, 8, 17, 24, 38, 63]. One shared intuition among such methods is that they utilize forgery augmentations to imitate the deepfake generation pipeline to encourages detection models to learn generic representative features.

Despite the success of such forgery augmentation-based methods, most of them exploit p-fake samples for training models in only two ways: 1) utilizing solely p-fake samples without the incorporation of o-fake samples, or 2) creating some p-fake samples and then mixing them into o-fake samples. In other words, the number of p-fake samples and the policy of forgery augmentation are fixed when training the deepfake detector. It may lead to inefficient training for the following reasons: First, applying forgery augmentation does not always bring improvement over the whole process of training. For instance, we observed that a detection model tends to learn faster during earlier training stages without using forgery augmentation. We hypothesize that models at the early stage of training still lacks the capability to recognize the original forgeries, so excessively introduced p-fake samples at such stages are not conducive to the convergence of the models. Secondly, using only a single type of forgery augmentation scheme to generate p-fake samples during the training is not optimal for the model. Intuitively, the detection model can learn more clues from p-fake samples synthesized by diverse kinds of forgery augmentation operations. Moreover, the




Fig. 1: The proposed CDFA adjusts the composition of fake samples during the training by: 1) gradually increasing the proportion of p-fake samples, and 2) applying a dynamic forgery augmentation policy to generate p-fake samples.

optimal forgery augmentation scheme should be different for every sample on variation of training stages.

Motivated by aforementioned concern, in this work, we propose a novel Curricular Dynamic Forgery Augmentation (CDFA) strategy. CDFA is a simple yet efficient method to improve the generalization for deepfake detectors by adjusting the composition of fake samples at different training stages (see Figure 1). As for the number of p-fake samples, we design a *Monotonic Curriculum (MC)* strategy that progressively introduces more p-fake samples while reducing the o-fake samples as training proceeds. Although the monotonic curriculum gradually increases the p-fake samples as the model improves, it does not determine which forgery augmentation operation applied to each sample can bring the most improvement to the model training. Motivated by the automatic augmentation paradigm [10, 12], we propose a *Dynamic Forgery Search (DFS)* strategy which considers the evaluation of the current model on the validation set as an expert to guide the optimization of which forgery augmentation operation is applied to each sample in different training stages. Furthermore, considering the current forgery augmentations [27, 45] can not imitate the temporal inconsistency of deepfake generation, we propose a novel forgery augmentation named *Self-shifted Blending Image (SSBI)*. It can simply introduce the temporal artifacts by blending the faces of two different frames from the same video. Comprehensive experimental results show that our method can significantly improve the generalization performances of various naive deepfake detectors in a plug-and-play manner, make them achieve superior performances over several SOTA competitors in multiple cross-datasets and cross-manipulations benchmarks.

Briefly, the main contributions of this work can be summarized as follows:

- To the best of our knowledge, it is the first work to investigate the p-fake sample scheme, including its proportion and generation method, during the training of deepfake detector.
- We propose a monotonic curriculum strategy that gradually introduces the proportion of p-fake samples along with the training process. Through such

easy-to-hard data strategy, we can improve the generalization performance while accelerating convergence of the deepfake detector.

- We propose a dynamic forgery search strategy that trains a policy network on the fly with the training of deepfake detector, which aims to search a optimal forgery augmentation policy based on evolving data and model states in different training stages.
- We further propose a novel forgery augmentation method, named Self-shifted Blending Image (SSBI), to compensate the deficiency of prior works in simulating temporal artifacts.

2 Related Works

Deepfake detection. The past five years have witnessed a wide variety of methods proposed for defending against the malicious usage of deepfakes. Early works focus on hand-crafted features such as eyes-blinking [28], inconsistent head poses [61] and visual artifacts [30, 36]. By formulating the detecting as a vanilla binary classification problem (i.e. pristine or forgery), current end-to-end trained detectors [1, 11, 50] to directly distinguish deepfake content from authentic data. To this end, several works [25, 32, 42] utilize frequency information to improve the performance of detectors. Moreover, there are some works aiming to localize the forged regions and make a decision based on the predicted regions [6]. Due to the development of deep generative models, the forged faces become more realistic and the manipulation methods are of more diversity. Some works propose to find clues on inconsistency of facial identity [13, 16, 17, 41]. Several works show introduce common data augmentations (e.g. blurring and jpeg compression) [4, 51, 60] can help improve the detection performance. Furthermore, [37] proposes to use RL agent to search the policy of common data augmentations (e.g., Brightness and Contrast). However, the improvement in generalization performance of the commonly data augmentation is limited.

Deepfake detection through forgery augmentation. One of the most effective approaches to improve generalization performance is to introduce forgery augmentation techniques to first synthesize forged images (i.e., pseudo deepfakes) and then train a deepfake detector model. As the pioneering works, BI [27] are introduced to generate blended faces which reproduce blending artifacts from pairs of two pristine images with similar facial landmarks. Following that, SBI [45] selects two views of the same face image as the target face and the source face. SLADD [7] employs an adversarial training strategy to find the most difficult BI configuration and trained a classifier to predict the forgeries. Recent works [2, 8, 17, 24, 38, 63] also conduct such forgery augmentation paradigm as a core part of improving generalization performances.

3 Methodology

In this section, we describe the proposed CDFA method in detail. The pipeline of CDFA is shown in Figure 2. First, we propose a monotonic curriculum strategy,

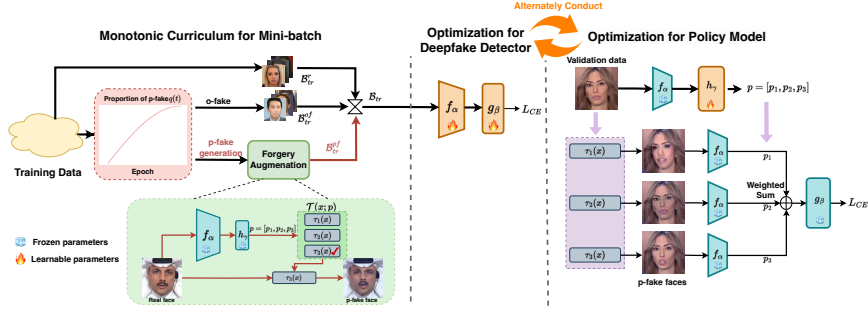


Fig. 2: Overview of the proposed CDFA.

designed to gradually increase the proportion of p-fake samples and decrease the proportion of o-fake samples in each mini-batch as the training proceeds. As for p-fake generation, we propose a dynamic forgery search strategy that optimizes a lightweight policy network to determine the preferred forgery augmentation operation for producing p-fake samples in different training stages.

3.1 Monotonic Curriculum

Previous works [27, 45] of forgery augmentations are used to simply mix p-fake samples into o-fake samples or use them directly (without o-fake samples) and then conduct the model training. Herein, we conducted a simple study by training the model by using o-fake (i.e., baseline) or p-fake samples alone. Figure 3 shows that the model converges much slower when only trained with p-fake samples (generated by SBI [45]). This suggests that the model can not even recognize the original forgery artifacts at the very early stage of the training. In other words, for the deepfake detection task, the o-fake samples can be considered as easy samples, while the p-fake is more difficult. Consequently, introducing a large number of p-fake samples at the initial stage appears to be not optimal for achieving efficient model convergence.

Inspired by curriculum learning paradigm [3, 22, 47, 48, 53], we propose a easy-to-hard data strategy that adjust the proportion of p-fake samples along with the training process. We introduce the curriculum schedule $q(t)$ about the proportion of p-fake samples as follows:

$$q(t) = \sin(t/\epsilon) \quad (1)$$




Fig. 3: Validation loss on FF++ during the training.

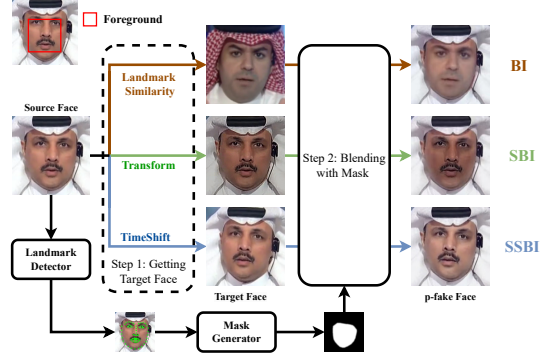


Fig. 4: Overall pipeline of forgery augmentations.

where t is the current training epoch number and ϵ is a manually adjustable hyper-parameter. We set $\epsilon = 2T/\pi$ to make $q(t)$ increases monotonically in $[0, 1]$, where T is the total number of the training epoch.

Let \mathcal{D}_{tr}^r and \mathcal{D}_{tr}^f be the real and fake part of the training data \mathcal{D}_{tr} , respectively. In constructing a training mini-batch \mathcal{B}_{tr} with batch size b , we first sample $b/2$ images from \mathcal{D}_{tr}^r as the real part, denoted as \mathcal{B}_{tr}^r . For the fake part of \mathcal{B}_{tr} , we compute number of o-fake and p-fake samples (denoted as n_{pf} and n_{of} respectively) by:

$$n_{pf} = q(t) \times b/2, n_{of} = (1 - q(t)) \times b/2 \quad (2)$$

For o-fake samples, \mathcal{B}_{tr}^{of} , we sample n_{of} images from \mathcal{D}_{tr}^f . For p-fake samples \mathcal{B}_{tr}^{pf} , we random select n_{pf} images from \mathcal{D}_{tr}^r and conduct forgery augmentations to generate them. Thus, \mathcal{B}_{tr} is obtained by:

$$\mathcal{B}_{tr} = \mathcal{B}_{tr}^r \cup \mathcal{B}_{tr}^{of} \cup \mathcal{B}_{tr}^{pf}, |\mathcal{B}_{tr}| = b \quad (3)$$

Utilizing this strategy, the model is primarily trained on the o-fake samples at the early stages, which facilitates rapid convergence by learning the obvious forgery traces in o-fake samples. As training proceeds, the model fully learns the original forgery artifacts and its training can benefit more from the augmented p-fake samples. To verify this, we observe that the convergence efficiency of the model (see in Figure 3) becomes higher after introducing the MC strategy.

3.2 Forgery Augmentation Operations

Given a pristine source face image x , the forgery augmentation operations can simply be considered as modifying the foreground face region while keeping the background. To achieve this, forgery augmentation generally consists of the two steps, i.e., 1) get a target face x_t , 2) blending it with a mask M . To generate the blending mask M , we first extract the facial landmarks $l(x)$ by Dlib [23]

Algorithm 1 Policy-Controlled Forgery Augmentation \mathcal{T}

Require: Source face $x \sim \mathcal{D}_{tr}^r$, augmentation policy p
Ensure: Pseudo-fake face \hat{x}

- 1: Get landmarks $l(x)$
- 2: Sample one operation j based on p
 $\quad \#$ *Getting the target face x_t*
- 3: **if** $j=1$ **then**
- 4: Conduct BI [27]: $x_t = \operatorname{argmin}_{x_t \sim \mathcal{D}_{tr}^r} |l(x_t) - l(x)|$
- 5: **else if** $j=2$ **then**
- 6: Conduct SBI [45]: $x_t = \operatorname{Transform}(x)$
- 7: **else if** $j=3$ **then**
- 8: Conduct SSBI: $x_t = \operatorname{TimeShift}(x, \operatorname{rand}(5, 10))$
- 9: **end if**
 $\quad \#$ *Blending with mask*
- 10: $M = \operatorname{Deform}(\operatorname{ConvexHull}(l(x)))$
- 11: Generate \hat{x} by Equation (4)

and then apply random deformation and blurring on the convex hull, which is inspired by [27, 45]. We obtain p-fake face \hat{x} by:

$$\hat{x} = x_t \odot M + x \odot (1 - M) \quad (4)$$

where \odot specifies the element-wise multiplication.

For the selection of x_t , BI [27], the target face is get from different identities with top facial similarity to the source face. As for SBI [45], the target face is get by source face itself with data transforms. However, the aforementioned works are dedicated to simulate the inconsistency in the spatial domain and thus cannot capture temporal inconsistencies across video frames, which is one of the important clues for identify deepfakes.

Self-shifted Blending Image: We propose a novel forgery augmentation operation named *Self-shifted Blending Image (SSBI)* to imitate temporal artifacts. The target face of SSBI is get from another frame on the same video. It can simply imitate temporal inconsistency between the foreground face and background in terms of face movements [61].

The overall pipeline of the aforementioned forgery augmentations can be seen in Figure 4.

3.3 Dynamic Forgery Search

Although the p-fake technology has been employed in some SOTA deepfake detection methods, the production of p-fake samples during the entire training period typically involves only a single forgery augmentation operation. Taking inspiration from the success of automatic data augmentation techniques [9, 10, 12, 29], we suggest that employing multiple forgery augmentation operations and dynamically adjusting their policy throughout the training process

is better than relying solely on a fixed single forgery augmentation operation for training the general deepfake detector. Thus, we propose to devise a strategy that dynamically selects optimal forgery augmentations based on evolving data and model states in different training stages.

To achieve this, we first define a policy-controlled forgery augmentation operator $\mathcal{T}(\cdot)$. Let \mathbb{T} be a set of forgery augmentation operations where τ_j denotes the j -th operation. In this work, \mathbb{T} only contains three forgery augmentation operations mentioned in Section 3.2 so that $|\mathbb{T}| = 3$. We formulate an forgery augmentation policy as the probability p in applying multiple forgery augmentation operations. Here, p is a probability vector with each entry: $p_j \in [0, 1]$; $\sum_{j=1}^{|\mathbb{T}|} p_j = 1$. As detailed in Algorithm 1, given an real face image x , we sample one operation according to an policy p and get the p-fake sample by

$$\hat{x} = \mathcal{T}(x; p) = \tau_j(x); j \sim p \quad (5)$$

The generated p-fake sample \hat{x} is labeled as fake.

Subsequently, we introduce the joint optimization of the deepfake detector and p-fake generation policy during the training. We employ a feature extraction network $f_\alpha : \mathcal{D} \rightarrow \mathcal{Z}$ to map a data space to a latent space, a classification head $g_\beta : \mathcal{Z} \rightarrow \mathcal{Y}$ to map a latent space to a label space. $g_\beta \circ f_\alpha$ can be regarded as a universal deepfake detector, where \circ is the compositional operator. We add a lightweight policy model $h_\gamma : \mathcal{Z} \rightarrow P$ to map a latent space to a probability space, where $p \in P$. The deepfake detector is optimized by minimizing the binary cross-entropy loss L_{CE} on a training batch \mathcal{B}_{tr} . The policy model is to search forgery augmentation policies applied to the training of the deepfake detector. Its optimization objective is to minimize L_{CE} on search batch data, denoted as \mathcal{B}_{sc} . Herein, we sample the \mathcal{B}_{sc} from the validation set \mathcal{D}_{val} .

Overall, the above objection can be formulated as a bi-level optimization problem [9, 33] as follow:

$$\begin{aligned} & \min_{\alpha, \beta} \mathcal{L}_{CE}(\alpha, \beta, \gamma^*; \mathcal{B}_{tr}) \\ & s.t. \min_{\gamma} \mathcal{L}_{CE}(\alpha^*, \beta^*, \gamma; \mathcal{B}_{sc}) \end{aligned} \quad (6)$$

We solve it by executing the following optimization phase alternatively.

Optimization for Deepfake Detector. In this phase, given the training data x , the frozen policy model h_γ generates the policy p . We get the augmented p-fake sample \hat{x} using Equation (5) and use them to train the detector model $g_\beta \circ f_\alpha$ by minimizing L_{CE} on the processed mini-batch \mathcal{B}_{tr} .

Optimization for Policy Model. In this phase, the weights of deepfake detector are frozen, and we aim to optimize h_γ policy given the validation data. However, we can not directly use back-propagation to optimize γ because the sampling process of one forgery augmentation operation in $\mathcal{T}(x; p)$ is non-differentiable. Hence, back-propagation cannot compute the partial derivative w.r.t. the augmentation probability p . To address this problem, we relax the non-differentiable $\mathcal{T}(x; p)$ to be a differentiable operator, The augmented validation data are passed

to the feature extraction network f individually to get the latent representations, which are then summed based on their weights in the probability vector p . The forward pass can be relaxed as the mixed representation and passed to g for computing the predicted labels:

$$\hat{y} = g_{\beta} \left(\sum_{j=1}^{|\mathbb{T}|} p_j \cdot f_{\alpha}(\tau_j(x)) \right); p = h_{\gamma} \circ f_{\alpha}(x) \quad (7)$$

In this way, we can update γ by minimizing L_{CE} combined with back-propagation.

With the aforementioned updating rules, both the policy and detector models can be alternatively optimized. We set the search frequency s to make that Optimization for Policy Model is executed after every s steps of Optimization for Deepfake Detector.

4 Experiments

4.1 Experiment Settings

Datasets and pre-processing. Following most previous works, we mainly conducted experiments on the FaceForensics++ (FF++) [43] dataset. It contains 1000 Pristine (PT) videos (i.e., the real sample) and 5000 fake videos forged by five manipulation methods, i.e., Deepfakes (DF), Face2Face (F2F), FaceSwap (FS), NeuralTextures (NT) and FaceShifter (FSh). Besides, FF++ provides three quality levels in compression for these videos: raw, high-quality (HQ) and low-quality (LQ). The **HQ version of FF++** is adopted by default in this paper. If any deviation from this default, it will be explicitly stated. The samples were split into disjoint training, validation, and testing sets at the video level follows the official protocol.

To demonstrate the performance of CDFA in cross-dataset settings, four additional datasets are adopted, i.e., Celeb-DF-v2 (CDF) [31], DeepFake Detection Challenge preview (DFDCP) [15] and DeepFake Detection Challenge public (DFDC) [14] and WildDeepfake(Wild) [65]. See the supplementary material for more details.

Implementation details. We use SwinTransformerV2-Base (Swin) [34] as the backbone network f_{α} , and the parameters are initialized by the weights pre-trained on the ImageNet. We implemented h_{γ} using three MLP layers with random initialization, and the softmax operation is applied to the output of h_{γ} to get the probabilities. We use the Adam optimizer for both the two networks with a cosine learning rate scheduler initiate with 0.0001. We set the total training epoch $T = 50$ and the searching frequency as $s = 10$. See the supplementary material for more details.

Evaluation Metrics. In this work, we mainly report the area under the ROC curve (AUC) to compare with prior works. The video-level results are obtained by averaging predictions over each frame on an evaluated video.

Table 1: Video-level (top) and frame-level (bottom) AUC(%) of cross-datasets performances compared with SOTA methods. The best results are highlighted.

Method	Backbone	Data	CDF	Wild	DFDCP	DFDC
TALL [56]	Swin	HQ	90.79	-	-	76.78
SeeABLE [24]	ENb4	HQ	87.30	-	86.30	75.90
CADDM [16]	ENb4	HQ	93.88	-	-	73.85
AUNet [2]	Xcep	HQ	92.77	-	86.16	73.82
LTTD [21]	Designed	HQ	89.30	-	-	80.40
CD-NET [46]	Xcep	HQ	88.50	-	-	77.00
DCL [49]	ENb4	HQ	82.30	71.14	-	76.71
PCL+I2G [63]	Res34	HQ	90.03	-	74.37	67.52
Ours	Swin	HQ	97.22	84.45	97.03	83.84
Ours	Swin	LQ	94.63	84.05	96.60	81.16
LSDA [58]	ENb4	HQ	83.00	-	81.50	73.60
UCF [59]	Xcep	HQ	75.27	-	75.94	71.91
SFDG [54]	ENb4	LQ	75.83	69.27	-	73.64
NoiseDF [52]	Designed	HQ	75.89	-	-	63.89
OST [8]	Xcep	HQ	74.80	-	-	83.30
UIA-ViT [64]	ViT-B	HQ	82.41	-	75.80	-
SLADD [7]	Xcep	HQ	79.70	-	-	77.20
RECCE [5]	Xcep	LQ	68.71	64.31	-	69.06
PEL [19]	ENb4	LQ	69.18	67.39	-	63.31
Ours	Swin	HQ	91.96	81.34	93.30	81.45
Ours	Swin	LQ	89.88	80.99	92.65	78.67

4.2 Generalization Comparisons

To comprehensively evaluate the generalizability of our method, we compare the performances of cross-datasets and cross-manipulation evaluations with several SOTA methods published in the past three years.

Cross-datasets evaluations. The cross-datasets evaluation is still a challenging task because the unknown domain gap between the training and testing datasets can be caused by different source data, forgery methods, and/or post-processing. In this part, we evaluate the generalization performances in a cross-dataset setting. Specifically, our models were trained on the FF++ (only containing DF, F2F, FS, and NT subsets for fair comparisons) and tested on other datasets. The experimental results in terms of frame-level and video-level AUC are shown in Table 1. We can observe that our method outperforms the best competition in terms of video-level evaluations. For frame-level evaluations, our method still outperforms most of the SOTA competitors regardless it is trained on the HQ or LQ version of FF++. For instance, our approach surpasses TALL [56], which also employs Swin as its backbone network, by around

Table 2: Video-level AUC(%) on cross-dataset evaluations (trained on FF++) performances. The best results are highlighted.

Methods	CDF	Wild	DFDCP	DFDC	Avg
Xcep	69.99	60.06	80.93	65.85	69.21
+ CDFA	93.24	78.06	86.43	77.06	84.56
ENb4	74.50	61.45	82.28	68.15	71.60
+ CDFA	95.81	78.76	86.66	78.38	85.25
Swin	73.53	71.56	89.37	71.30	76.44
+ CDFA	97.22	84.45	97.03	83.84	91.63

Table 3: Video-level AUC(%) on cross-manipulation evaluations. Trained on FF++/DF. The best cross-manipulation results are highlighted.

Method	DF	F2F	FS	NT	FSh
CADDM [16]	100	83.94	58.33	68.98	-
DCL [49]	99.98	77.13	61.01	75.01	-
Xcep	100	73.60	53.73	71.53	71.62
+ CDFA	99.54	85.93	90.04	82.23	77.81
ENb4	100	71.23	47.32	70.31	75.71
+ CDFA	99.65	88.51	91.84	82.14	84.57
Swin	100	67.43	56.74	78.74	70.87
+ CDFA	99.90	87.44	90.64	86.27	76.64

7% when testing on CDF and DFDC. We can also see that our method obtains a lower frame-level AUC when testing on DFDC compared to OST [8]. One possible explanation is that OST introduces a test-time adaptation strategy that adapts the model with domain knowledge of testing data before evaluation. This trick facilitates for evaluating large-scale unseen data such as DFDC. However, our method never introduces knowledge from testing data during the training.

Backbone impact. In Table 2, we evaluated the performances of CDFA when employing different backbone architectures f_α , i.e., Xception (Xcep) [11], EfficientNetb4 (ENb4) [50] and Swin [34]. We observe that our CDFA can significantly improve the generalization performances of all evaluated models (at least 13% on average). In conjunction with the results in Table 1, CDFA still outperforms the SOTA competitors even with the same backbones (e.g., the comparison of SeeABLE and ENb4+CDFA). We also find that larger and more powerful encoders lead to better generalization in general when equipping CDFA. These results suggest that CDFA is applicable to different backbone models and is expected to further benefit from future developments in model topologies.

Cross manipulation evaluations. In real detection situations, the defenders generally are not aware of the attacker’s forgery methods. For this reason, it is important to verify the model generalization to various forgery methods. We conducted the cross-manipulation experiment on FF++, all models were trained on the DF subset and tested on the remaining four manipulations. More results on other subsets are given in the Supplementary Material. We evaluate the effect of different backbone architectures f_α equipped with the proposed CDFA. As shown in Table 3, we can observe that our CDFA can improve cross-manipulation performances significantly regardless of the types of backbones. In addition, the backbone models trained with the CDFA approach outperform the SOTA competitors (i.e, CADDM [16] and DCL [49]) by a considerable margin on average. These results highlight the effectiveness of CDFA in combating emerging unseen forgery methods.

Table 4: Video-level AUC(%) performances for ablation studies.

Variant	Fake data		Strategies		CDF	Wild	DFDCP	DFDC	Avg
	o-fake	p-fake	MC	DFS					
1	-	BI	-	-	88.40	82.71	87.86	74.55	83.38
2	-	SBI	-	-	91.31	76.14	91.57	70.14	82.29
3	-	SSBI	-	-	90.60	83.31	95.98	76.76	86.67
4	-	ALL	-	-	95.03	77.39	90.49	77.41	85.08
5	-	ALL	-	✓	95.84	82.47	93.18	82.27	88.44
6	✓	BI	✓	-	93.34	80.16	89.64	78.37	85.38
7	✓	SBI	✓	-	94.41	75.37	88.36	73.54	82.92
8	✓	SSBI	✓	-	94.75	83.30	95.52	83.95	89.38
9	✓	ALL	✓	-	97.26	85.22	96.46	82.20	90.29
10	✓	ALL	-	-	96.31	81.57	95.61	83.79	89.32
11	✓	ALL	-	✓	97.06	84.72	95.73	86.37	90.97
CDFA	✓	ALL	✓	✓	97.94	86.72	97.63	87.16	92.36

4.3 Ablation Studies

In this part, we perform several ablations to better understand the contributions of each component in the proposed CDFA, including consists of fake data during the training, monotonic curriculum and dynamic forgery search. We evaluated several variants of the proposed CDFA (Trained on FF++/DF) and summarized the results in Table 4.

Effects of the forgery augmentation operations. From the comparison among *Variant 1*, *Variant 2*, and *Variant 3*, where the fake part of the training data only contained p-fake samples generated by single forgery augmentation operation, we observe that our proposed SSBI performs better compared to BI and SBI in most results. Similar phenomena also appear in the comparison of *Variant 6*, *7*, and *Variant 8*. These results highlight the effectiveness of the proposed SSBI which compensates the deficiency in simulating temporal artifacts.

Effects of monotonic curriculum. From the comparative analysis of *Variant 1* with *Variant 6*, *Variant 2* with *Variant 7*, *Variant 3* with *Variant 8*, and *Variant 4* with *Variant 9*, we note that using of o-fake samples through MC strategies can improve the performances in most cases compared to only training with p-fake samples. Furthermore, the comparison between *Variant 9* and *Variant 10* indicates that when using o-fake, using MC strategies performs better than simply mixing them into p-fake samples. It shows the effectiveness of the idea behind MC, which is to use o-fake at the beginning and gradually introduce more p-fake during the training.

Effects of dynamic forgery search. The comparative analysis of the *Variant 3*, *Variant 4* and *Variant 5* indicates that utilizing all forgery augmentation operations with fixed probability can not improve the performances compared with using SSBI alone, while the introduction of DFS leads to significant improvements. Furthermore, from the comparative analysis of the *Variant 5*, *Variant 11*

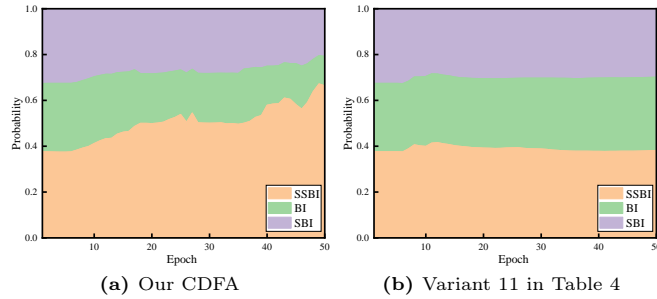


Fig. 5: Evolution of policies searched by DFS for two variants method: **(a)** the proposed CDFA, **(b)** Variant 11 in Table 4. The policy of one epoch is obtained by averaging the policies searched at all steps over the epoch.

and CDFA, we can find that the improvements brought by DFS can be further increased with the guidance of MC. It shows the effectiveness of DFS which dynamically searches the optimal forgery augmentation policy during the training.

4.4 Visualizations and Analysis

Analysis of searched policies. In this part, we depict and analyze the evolution of policies searched by DFS throughout the training process. For CDFA, as shown in Figure 5a, we observe a progressive rise in the probability of SSBI as training proceeds, whereas the probabilities of BI and SBI gradually decline. Such a phenomenon suggests that the model places increasing importance on p-fake samples generated by SSBI as the training proceeds. This observed evolution of policies also aligns with the results in Table 4, i.e., the performances of using SSBI alone surpass that of BI and SBI alone. It further emphasizes the effectiveness of the proposed SSBI. Moreover, from Figure 5b, it is apparent that maintaining a constant proportion of o-fake samples throughout the training does not leverage the full potential of DFS. This constancy may cause DFS to become less dynamic, failing to adaptively adjust the training strategy to optimize deepfake detection effectively. These results highlight the importance of the guidance introduced by MC strategy.

Analysis of fake samples in training process. In this part, we study the properties of fake samples during the training. Specifically, we employed a baseline deepfake detector (i.e., Swin [34]) as an assessment model. We first train the assessment model on FF++ and then fix it to assess the fake samples utilized in each training epoch. We depict the assessing accuracy of our CDFA and fixed policy of o-fake and p-fake samples (i.e., *Variant 10* in Table 4) in Figure 6. It can be observed that assessment accuracy decreases significantly in the early stage of training, while it fluctuates in the later stages. This observation suggests that our CDFA gradually increases the difficulty of fake samples via MC in the early stages of training while maintaining their diversity via DFS in the later stages of training when p-fake dominates the fake samples. It reveals that the




Fig. 6: Assessing accuracy during the training.

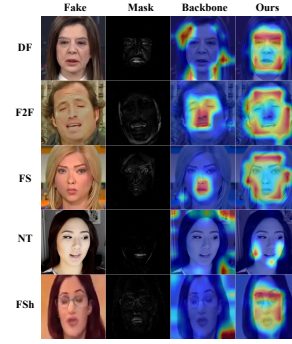


Fig. 7: GradCAM visualizations of the backbone and our proposed CDFA.

deepfake detector can learn more general forgery representations by gradually focusing on hard fake samples with diversity during the training.

Visualizations of CAM. To intuitively demonstrate which patterns learned with our proposed CDFA, we compare the GradCAM [44] visualizations between the backbone model (i.e., Swin) trained with and without CDFA. As shown in Figure 7, the backbone trained without CDFA tend to only capture method-specific artifacts. Our method identifies the forgery faces by focusing on the general artifacts (e.g., the blending traces on the boundary between the background and foreground face) with the help of CDFA. Based on the results of our quantitative experiments (Table 2), we believe that paying attention to the inconsistency between the background and facial parts can improve the generalization ability of deepfake detectors.

5 Conclusions

In this work, we present CDFA, which aims to improve the generalization performances of deepfake detectors by dynamically adjusting the composition of fake samples during the training. First, we propose a monotonic curriculum that progressively increases the proportion of p-fake samples as training proceeds. Second, we propose a dynamic forgery searching strategy to conduct the optimal forgery augmentation operation for each image varying between training stages. In addition, we propose a novel forgery augmentation scheme named SSBI to simply imitate the temporal inconsistency of deepfake generation. Comprehensive experiments show that CDFA can significantly improve the performances of various naive deepfake detectors in a plug-and-play way, and make them attain superior performances over the existing methods in several cross-datasets and cross-manipulations benchmarks.

Ethic Statement. All face images used in this paper were obtained from public datasets. There is no violation of personal privacy while conducting experiments in this work.

Acknowledgments

This work was supported in part by National Natural Science Foundation of China (Grant U23B2022, U22B2047, U22A2030), Guangdong Provincial Key Laboratory (Grant 2023B1212060076) and Guangdong Major Project of Basic and Applied Basic Research (Grant No. 2023B0303000010). The work was also supported in part by China Postdoctoral Science Foundation under Grant 2021TQ0314 and Grant 2021M703036.

References

1. Afchar, D., Nozick, V., Yamagishi, J., Echizen, I.: MesoNet: A Compact Facial Video Forgery Detection Network. In: 2018 IEEE International Workshop on Information Forensics and Security (WIFS). pp. 1–7 (2018)
2. Bai, W., Liu, Y., Zhang, Z., Li, B., Hu, W.: AUNet: Learning Relations Between Action Units for Face Forgery Detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 24709–24719 (2023)
3. Bengio, Y., Louradour, J., Collobert, R., Weston, J.: Curriculum learning. In: Proceedings of the 26th Annual International Conference on Machine Learning. pp. 41–48 (2009)
4. Bondi, L., Cannas, E.D., Bestagini, P., Tubaro, S.: Training Strategies and Data Augmentations in CNN-based DeepFake Video Detection. In: 2020 IEEE International Workshop on Information Forensics and Security (WIFS). pp. 1–6 (2020)
5. Cao, J., Ma, C., Yao, T., Chen, S., Ding, S., Yang, X.: End-to-End Reconstruction-Classification Learning for Face Forgery Detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4113–4122 (2022)
6. Chen, H., Lin, Y., Li, B.: Exposing Face Forgery Clues via Retinex-based Image Enhancement. In: Proceedings of the Asian Conference on Computer Vision. pp. 602–617 (2022)
7. Chen, L., Zhang, Y., Song, Y., Liu, L., Wang, J.: 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. pp. 18710–18719 (2022)
8. Chen, L., Zhang, Y., Song, Y., Wang, J., Liu, L.: OST: Improving Generalization of DeepFake Detection via One-Shot Test-Time Training. *Advances in Neural Information Processing Systems* **35**, 24597–24610 (2022)
9. Cheung, T.H., Yeung, D.Y.: AdaAug: Learning Class- and Instance-adaptive Data Augmentation Policies. In: International Conference on Learning Representations (2022)
10. Cheung, T.H., Yeung, D.Y.: A Survey of Automated Data Augmentation for Image Classification: Learning to Compose, Mix, and Generate. *IEEE Transactions on Neural Networks and Learning Systems* pp. 1–21 (2023)
11. Chollet, F.: Xception: Deep Learning With Depthwise Separable Convolutions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1251–1258 (2017)
12. Cubuk, E.D., Zoph, B., Mane, D., Vasudevan, V., Le, Q.V.: AutoAugment: Learning Augmentation Strategies From Data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 113–123 (2019)

13. Das, S., Kolahdouzi, M., Özparlak, L., Hickie, W., Etemad, A.: Unmasking Deepfakes: Masked Autoencoding Spatiotemporal Transformers for Enhanced Video Forgery Detection. In: 2023 IEEE International Joint Conference on Biometrics (IJCB). pp. 1–11 (2023)
14. Dolhansky, B., Bitton, J., Pflaum, B., Lu, J., Howes, R., Wang, M., Ferrer, C.C.: The DeepFake Detection Challenge (DFDC) Dataset. arXiv:2006.07397 [cs] (2020)
15. Dolhansky, B., Howes, R., Pflaum, B., Baram, N., Ferrer, C.C.: The Deepfake Detection Challenge (DFDC) Preview Dataset. arXiv:1910.08854 [cs] (2019)
16. Dong, S., Wang, J., Ji, R., Liang, J., Fan, H., Ge, Z.: Implicit Identity Leakage: The Stumbling Block to Improving Deepfake Detection Generalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3994–4004 (2023)
17. Dong, X., Bao, J., Chen, D., Zhang, T., Zhang, W., Yu, N., Chen, D., Wen, F., Guo, B.: Protecting Celebrities From DeepFake With Identity Consistency Transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9468–9478 (2022)
18. Gao, G., Huang, H., Fu, C., Li, Z., He, R.: Information Bottleneck Disentanglement for Identity Swapping. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3404–3413 (2021)
19. Gu, Q., Chen, S., Yao, T., Chen, Y., Ding, S., Yi, R.: Exploiting Fine-Grained Face Forgery Clues via Progressive Enhancement Learning. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 735–743 (2022)
20. Guan, J., Zhou, H., Gong, M., Ding, E., Wang, J., Zhao, Y.: Detecting Deepfake by Creating Spatio-Temporal Regularity Disruption (2023)
21. Guan, J., Zhou, H., Hong, Z., Ding, E., Wang, J., Quan, C., Zhao, Y.: Delving into Sequential Patches for Deepfake Detection. *Advances in Neural Information Processing Systems* **35**, 4517–4530 (2022)
22. Hou, C., Zhang, J., Zhou, T.: When to Learn What: Model-Adaptive Data Augmentation Curriculum. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1717–1728 (2023)
23. King, D.E.: Dlib-ml: A Machine Learning Toolkit. *The Journal of Machine Learning Research* **10**, 1755–1758 (2009)
24. Larue, N., Vu, N.S., Struc, V., Peer, P., Christophides, V.: SeeABLE: Soft Discrepancies and Bounded Contrastive Learning for Exposing Deepfakes. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 21011–21021 (2023)
25. Li, J., Xie, H., Li, J., Wang, Z., Zhang, Y.: Frequency-Aware Discriminative Feature Learning Supervised by Single-Center Loss for Face Forgery Detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6458–6467 (2021)
26. Li, L., Bao, J., Yang, H., Chen, D., Wen, F.: Advancing High Fidelity Identity Swapping for Forgery Detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5074–5083 (2020)
27. Li, L., Bao, J., Zhang, T., Yang, H., Chen, D., Wen, F., Guo, B.: Face X-Ray for More General Face Forgery Detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5001–5010 (2020)
28. Li, Y., Chang, M., Lyu, S.: In Ictu Oculi: Exposing AI Created Fake Videos by Detecting Eye Blinking. In: 2018 IEEE International Workshop on Information Forensics and Security (WIFS). pp. 1–7 (2018)

29. Li, Y., Hu, G., Wang, Y., Hospedales, T., Robertson, N.M., Yang, Y.: Differentiable Automatic Data Augmentation. In: Computer Vision – ECCV 2020. pp. 580–595 (2020)
30. Li, Y., Lyu, S.: Exposing DeepFake Videos By Detecting Face Warping Artifacts. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 46–52 (2019)
31. Li, Y., Yang, X., Sun, P., Qi, H., Lyu, S.: Celeb-DF: A Large-Scale Challenging Dataset for DeepFake Forensics. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3207–3216 (2020)
32. Liu, H., Li, X., Zhou, W., Chen, Y., He, Y., Xue, H., Zhang, W., Yu, N.: Spatial-Phase Shallow Learning: Rethinking Face Forgery Detection in Frequency Domain. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 772–781 (2021)
33. Liu, R., Gao, J., Zhang, J., Meng, D., Lin, Z.: Investigating Bi-Level Optimization for Learning and Vision From a Unified Perspective: A Survey and Beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(12), 10045–10067 (2022)
34. Liu, Z., Hu, H., Lin, Y., Yao, Z., Xie, Z., Wei, Y., Ning, J., Cao, Y., Zhang, Z., Dong, L., Wei, F., Guo, B.: Swin Transformer V2: Scaling Up Capacity and Resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12009–12019 (2022)
35. Luo, Y., Zhang, Y., Yan, J., Liu, W.: Generalizing Face Forgery Detection With High-Frequency Features. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16317–16326 (2021)
36. Matern, F., Riess, C., Stamminger, M.: Exploiting Visual Artifacts to Expose Deepfakes and Face Manipulations. In: 2019 IEEE Winter Applications of Computer Vision Workshops (WACVW). pp. 83–92 (2019)
37. Nadimpalli, A.V., Rattani, A.: On Improving Cross-dataset Generalization of Deepfake Detectors. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 91–99 (2022)
38. Nguyen, D., Mejri, N., Singh, I.P., Kuleshova, P., Astrid, M., Kacem, A., Ghorbel, E., Aouada, D.: LAA-Net: Localized Artifact Attention Network for Quality-Agnostic and Generalizable Deepfake Detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17395–17405 (2024)
39. Nirkin, Y., Keller, Y., Hassner, T.: FSGAN: Subject Agnostic Face Swapping and Reenactment. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7184–7193 (2019)
40. Nirkin, Y., Keller, Y., Hassner, T.: FSGANv2: Improved Subject Agnostic Face Swapping and Reenactment. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(1), 560–575 (2023)
41. Nirkin, Y., Wolf, L., Keller, Y., Hassner, T.: DeepFake Detection Based on Discrepancies Between Faces and Their Context. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(10), 6111–6121 (2022)
42. Qian, Y., Yin, G., Sheng, L., Chen, Z., Shao, J.: Thinking in Frequency: Face Forgery Detection by Mining Frequency-Aware Clues. In: ECCV. pp. 86–103 (2020)
43. Rossler, A., Cozzolino, D., Verdoliva, L., Riess, C., Thies, J., Niessner, M.: Face-Forensics++: Learning to Detect Manipulated Facial Images. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1–11 (2019)
44. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-CAM: Visual Explanations From Deep Networks via Gradient-Based Localization.

- In: Proceedings of the IEEE International Conference on Computer Vision. pp. 618–626 (2017)
45. Shiohara, K., Yamasaki, T.: Detecting Deepfakes With Self-Blended Images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18720–18729 (2022)
 46. Song, L., Fang, Z., Li, X., Dong, X., Jin, Z., Chen, Y., Lyu, S.: Adaptive Face Forgery Detection in Cross Domain. In: Computer Vision – ECCV 2022. pp. 467–484 (2022)
 47. Song, W., Lin, Y., Li, B.: Towards Generic Deepfake Detection with Dynamic Curriculum. In: ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 4500–4504 (2024)
 48. Soviany, P., Ionescu, R.T., Rota, P., Sebe, N.: Curriculum Learning: A Survey. *International Journal of Computer Vision* **130**(6), 1526–1565 (2022)
 49. Sun, K., Yao, T., Chen, S., Ding, S., Li, J., Ji, R.: Dual Contrastive Learning for General Face Forgery Detection. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 2316–2324 (2022)
 50. Tan, M., Le, Q.: EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In: International Conference on Machine Learning. pp. 6105–6114. PMLR (2019)
 51. Wang, S.Y., Wang, O., Zhang, R., Owens, A., Efros, A.A.: CNN-Generated Images Are Surprisingly Easy to Spot... for Now. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8695–8704 (2020)
 52. Wang, T., Chow, K.P.: Noise Based Deepfake Detection via Multi-Head Relative-Interaction. *Proceedings of the AAAI Conference on Artificial Intelligence* **37**(12), 14548–14556 (2023)
 53. Wang, X., Chen, Y., Zhu, W.: A Survey on Curriculum Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(9), 4555–4576 (2022)
 54. Wang, Y., Yu, K., Chen, C., Hu, X., Peng, S.: Dynamic Graph Learning With Content-Guided Spatial-Frequency Relation Reasoning for Deepfake Detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7278–7287 (2023)
 55. Xu, C., Zhang, J., Han, Y., Tian, G., Zeng, X., Tai, Y., Wang, Y., Wang, C., Liu, Y.: Designing One Unified Framework for High-Fidelity Face Reenactment and Swapping. In: Computer Vision – ECCV 2022. pp. 54–71 (2022)
 56. Xu, Y., Liang, J., Jia, G., Yang, Z., Zhang, Y., He, R.: TALL: Thumbnail Layout for Deepfake Video Detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 22658–22668 (2023)
 57. Xu, Z., Zhou, H., Hong, Z., Liu, Z., Liu, J., Guo, Z., Han, J., Liu, J., Ding, E., Wang, J.: StyleSwap: Style-Based Generator Empowers Robust Face Swapping. In: Computer Vision – ECCV 2022. pp. 661–677 (2022)
 58. Yan, Z., Luo, Y., Lyu, S., Liu, Q., Wu, B.: Transcending Forgery Specificity with Latent Space Augmentation for Generalizable Deepfake Detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8984–8994 (2024)
 59. Yan, Z., Zhang, Y., Fan, Y., Wu, B.: UCF: Uncovering Common Features for Generalizable Deepfake Detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 22412–22423 (2023)
 60. Yan, Z., Zhang, Y., Yuan, X., Lyu, S., Wu, B.: DeepfakeBench: A Comprehensive Benchmark of Deepfake Detection. *Advances in Neural Information Processing Systems* **36**, 4534–4565 (Dec 2023)

61. Yang, X., Li, Y., Lyu, S.: Exposing Deep Fakes Using Inconsistent Head Poses. In: ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 8261–8265 (2019)
62. Zhao, H., Zhou, W., Chen, D., Wei, T., Zhang, W., Yu, N.: Multi-Attentional Deepfake Detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2185–2194 (2021)
63. Zhao, T., Xu, X., Xu, M., Ding, H., Xiong, Y., Xia, W.: Learning Self-Consistency for Deepfake Detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15023–15033 (2021)
64. Zhuang, W., Chu, Q., Tan, Z., Liu, Q., Yuan, H., Miao, C., Luo, Z., Yu, N.: UIA-ViT: Unsupervised Inconsistency-Aware Method Based on Vision Transformer for Face Forgery Detection. In: Computer Vision – ECCV 2022. pp. 391–407 (2022)
65. Zi, B., Chang, M., Chen, J., Ma, X., Jiang, Y.G.: WildDeepfake: A Challenging Real-World Dataset for Deepfake Detection. In: Proceedings of the 28th ACM International Conference on Multimedia. pp. 2382–2390 (2020)