

# DomainForensics: Exposing Face Forgery Across Domains via Bi-Directional Adaptation

Qingxuan Lv<sup>ID</sup>, Yuezun Li<sup>ID</sup>, *Member, IEEE*, Junyu Dong<sup>ID</sup>, Sheng Chen<sup>ID</sup>, *Life Fellow, IEEE*,  
Hui Yu<sup>ID</sup>, Huiyu Zhou<sup>ID</sup>, and Shu Zhang<sup>ID</sup>

**Abstract**—Recent DeepFake detection methods have shown excellent performance on public datasets but are significantly degraded on new forgeries. Solving this problem is important, as new forgeries emerge daily with the continuously evolving generative techniques. Many efforts have been made for this issue by seeking the commonly existing traces empirically on data level. In this paper, we rethink this problem and propose a new solution from the unsupervised domain adaptation perspective. Our solution, called *DomainForensics*, aims to transfer the forgery knowledge from known forgeries (fully labeled source domain) to new forgeries (label-free target domain). Unlike recent efforts, our solution does not focus on data view but on learning strategies of DeepFake detectors to capture the knowledge of new forgeries through the alignment of domain discrepancies. In particular, unlike the general domain adaptation methods which consider the knowledge transfer in the semantic class category, thus having limited application, our approach captures the subtle forgery traces. We describe a new bi-directional adaptation strategy dedicated to capturing the forgery knowledge across domains. Specifically, our strategy considers both forward and backward adaptation, to transfer the forgery knowledge from the source domain to the target domain in forward adaptation and then reverse the adaptation from the target domain to the source domain in backward adaptation. In forward adaptation, we perform supervised training for the DeepFake detector in the source domain and jointly employ adversarial feature adaptation to transfer the ability to detect manipulated faces from known forgeries to new forgeries. In backward adaptation, we further improve the knowledge transfer by coupling adversarial adaptation with self-distillation on new forgeries. This enables

the detector to expose new forgery features from unlabeled data and avoid forgetting the known knowledge of known forgery. Extensive experiments demonstrate that our method is surprisingly effective in exposing new forgeries, and can be plug-and-play on other DeepFake detection architectures.

**Index Terms**—Digital forensics, DeepFake detection, Domain-Forensics.

## I. INTRODUCTION

THE ever-growing convolutional neural network (CNN) based generative models [1], [2], [3], [4], [5], [6] have made face forgery much easier than ever before, allowing people to manipulate the face's identity, appearance and attributes in high realism with little effort. These CNN-based face forgery techniques, known as *DeepFake*, have drawn much attention, as their abuse using can lead to impersonation videos, economic fraud, biometric attacks, and even national security problems [7]. Thus, it is urgent and important to counteract the misuse of DeepFakes.

During the past few years, large number of DeepFake detection methods [8], [9], [10], [11], [12], [72], [73] have emerged. Trained on the recently proposed large DeepFake datasets, such as FaceForensics++ (FF++) [13] and Celeb-DF [14], these detection methods have shown promising performance. However, these methods fall into the category that the training and testing sets are from the same distribution, e.g., the same type of forgery or the same dataset, which unfortunately limits their practical applications, as there are always new types of forgeries emerging continuously and widespreading to everywhere on various social platforms. These new types of forgeries are very unlikely to have been included in the existing datasets, and thus they are unseen to these detectors, causing significant performance degradation (see Fig. 1 top part). This circumstance gives rise to a big challenge to DeepFake detectors, that is, how to detect constantly emerging new forgeries.

Recently, attempts have been made in the literature to solve this issue. One typical line of research is to use a variety of data augmentation to increase the generalization ability [15], [16], [17], [18]. These methods usually create forged faces by augmenting the pristine videos to cover the known types of forgeries as much as possible. Despite of the promisingly improved generalization, the types of augmentation are limited to known forgeries, thus hindering the performance when confronting unseen forgeries. Frequency clue is also used to improve generalization ability [19], [20], [21], [22].

Manuscript received 1 November 2023; revised 27 April 2024 and 2 July 2024; accepted 4 July 2024. Date of publication 18 July 2024; date of current version 2 August 2024. This work was supported in part by the National Key Research and Development Program of China under Grant 2022ZD0117201 and in part by the Sanya Science and Technology Special Fund under Grant 2022KJCX92. The work of Yuezun Li was supported by China Postdoctoral Science Foundation under Grant 2021TQ0314 and Grant 2021M703036. The associate editor coordinating the review of this article and approving it for publication was Dr. Benedetta Tondi. (*Corresponding authors: Yuezun Li; Junyu Dong.*)

Qingxuan Lv, Yuezun Li, Junyu Dong, and Shu Zhang are with the College of Computer Science and Technology, Ocean University of China, Qingdao 266100, China (e-mail: lvqingxuan@stu.ouc.edu.cn; liyuezun@ouc.edu.cn; dongjunyu@ouc.edu.cn; zhangshu@ouc.edu.cn).

Sheng Chen is with the School of Electronics and Computer Science, University of Southampton, SO17 1BJ Southampton, U.K., and also with the College of Computer Science and Technology, Ocean University of China, Qingdao 266100, China (e-mail: sqc@ecs.soton.ac.uk).

Hui Yu is with the School of Creative Technologies, Faculty of Creative and Cultural Industries, University of Portsmouth, PO1 2DJ Portsmouth, U.K. (e-mail: hui.yu@port.ac.uk).

Huiyu Zhou is with the School of Computing and Mathematical Sciences, University of Leicester, LE1 7RH Leicester, U.K. (e-mail: hz143@leicester.ac.uk).

Digital Object Identifier 10.1109/TIFS.2024.3426317

1556-6021 © 2024 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.  
See <https://www.ieee.org/publications/rights/index.html> for more information.

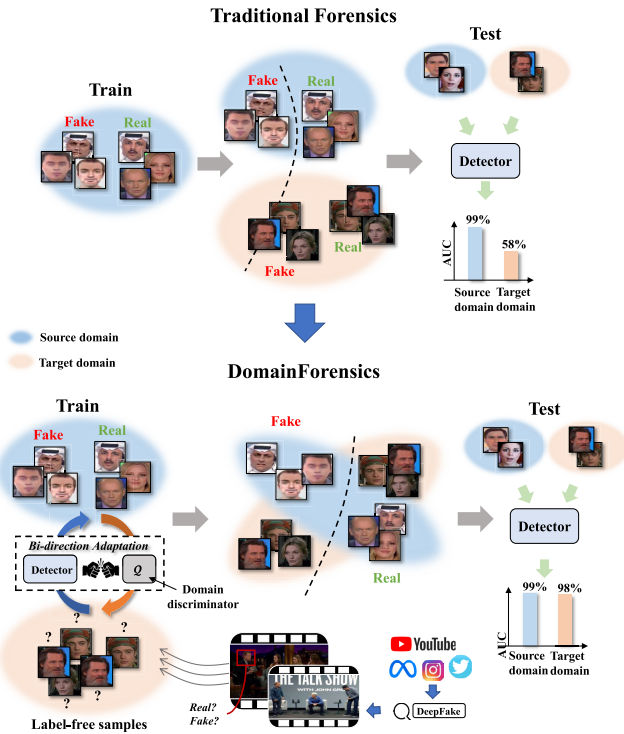


Fig. 1. Overview of traditional forensics (top) and *DomainForensics* (bottom). Traditional forensics achieves excellent performance on known forgeries but performs poorly on new forgeries. In contrast, *DomainForensics* can effectively expose new forgeries by performing the proposed bi-directional adaption, which can learn the common forgery features across domains using adversarial training.

However, this clue is easily affected by data processing and highly correlated with video quality, which cannot perform consistently across different datasets. A different direction of research is to apply transfer learning, such as zero- and few-shot learning [23], [24], [25], to improving generalization on new forgeries. Since zero-shot learning cannot access the samples of new forgeries in training, its performance is highly suppressed. In contrast, few-shot learning methods relax the restrictions in that they can access a few samples of new forgeries in training. However, this requires the annotation of these samples, which may not be easily obtained in practice, as we may not know whether a face is forged or not, e.g., multiple faces are in view but only video-level labels are provided. Thus, a fundamental question is: *can we detect new forgeries by only accessing target samples without any labels, while achieving competitive performance?*

In this paper, we cast DeepFake detection into a new formulation as an *unsupervised domain adaptation* problem, by transferring the knowledge from the source domain to the target domain, without using any annotations of target samples in training. This is very different from the existing strategies and it offers significant advantages over them. Specifically, for DeepFake detection, we can treat the known forgeries as the source domain and new forgeries as the target domain, see Fig. 1 bottom part. Our goal is to push the DeepFake detector to learn the common forgery features across different domains by only using label-free interested video collections. It is worth noting that this DeepFake detection problem has a significant discrepancy with the general unsupervised domain adaptation

problem, as we aim to learn the common forgery feature from the same category of faces (real or fake), which is more subtle than the semantic features of different categories in the general unsupervised domain adaptation problem (e.g., cat, dog, etc.). To this end, we propose a new unsupervised domain adaptation framework, called *DomainForensics*, for DeepFake detection. The key to our *DomainForensics* is a novel *bi-directional adaptation* strategy. This is very different from the existing DeepFake detection framework which only considers one direction to learn the knowledge supervised by the source domain and transfer it to the target domain. However, since the forgery features are subtle, the one-directional adaptation will inevitably lose a certain amount of knowledge [26], [27], [28], thus limiting the achievable performance on the target domain. To overcome this problem, we design bi-directional adaptation, which first transfers the knowledge from the source domain to the target domain, referred to as forward adaptation, and then reverses the adaptation from the target domain to the source domain, called backward adaptation. The backward adaptation stage utilizes the results of the forward adaptation stage, further explores the knowledge from the target domain, and transfers it back to the source domain. With the mutual adaptation, DeepFake detector can fully grab the common forgery features across domains.

To verify our idea, we adopt Vision Transformer (ViT) [29] as our DeepFake detector in the experiment, due to its successful application on vision tasks. Other architectures, such as ResNet [30], Xception [31] and EfficientNet [32], can also be used in our framework, and this will also be demonstrated. Since the frequency space can reveal the forgery traces [19], [20], we use color images and corresponding frequency-transformed maps as the input. In the forward adaptation stage, we develop a discriminator that is trained together with the DeepFake detector in an adversarial manner, where the discriminator aims to tell which domain the learned feature is from, and the DeepFake detector aims to extract features that confuse the discriminator. By doing so, the distribution of the target domain is pulled close to the source domain. In the backward adaptation stage, the adaptation is reverted. Since no labels are provided, we employ self-distillation [27] to further excavate the knowledge of the target domain, and then apply the adversarial training to the distilled model, in order to transfer the knowledge back to the source domain. Extensive experiments are conducted on FF++ and Celeb-DF datasets in several cross-domain scenarios, including different manipulation methods, datasets and types, to demonstrate the effectiveness of our method.

The contribution of this work is summarized as follows.

- 1) We propose a new DeepFake detection solution called *DomainForensics* to handle continuously emerged new forgeries. Different from recent efforts, our method focuses on pushing the detectors to learn the common forgery features across domains, that is, to transfer the forgery knowledge from known forgeries to unseen forgeries, instead of empirically blending faces on the data level.
- 2) We propose a new bi-directional adaptation strategy, which first transfers the forgery knowledge from the

source domain to the target domain in forward adaptation, and then reverses the adaptation from the target domain to the source domain in backward adaptation. Since the forgery traces are very subtle, we design the backward adaptation stage to further refine the results obtained from the forward adaptation stage with a self-distillation scheme.

- 3) Extensive experiments are conducted on FF++ and Celeb-DF datasets with several cross-domain scenarios, including crossing manipulation methods, crossing datasets, and crossing generative types, to demonstrate the effectiveness of our method. We also study the effects of various adaptation settings, various amounts of training samples and different components, to provide thoughtful insights for the following research.

The remainder of this paper is organized as follows. Section II reviews the recent works on DeepFake detection and unsupervised domain adaptation. Section III details our proposed *DomainForensics*, including the problem formulation, network framework and bi-directional adaptation. Section IV offers extensive experiments and elaborates on the experimental results. The paper concludes in Section V.

## II. RELATED WORK

In this section, we first present an overview of the existing deepfake detection approaches. We then provide a brief review of unsupervised domain adaptation and discuss the differences between the previous works and our approach.

### A. Deepfake Detection

With the advent of large-scale DeepFake datasets, e.g., [13] and [14], DeepFake detection has made significant progress in recent years, e.g., [8], [9], [10], [11], [12], [16], [18], [33], [34], [35], [36], [37], [72], and [73]. One challenging problem in this task is how to detect constantly emerging new forgeries. The methods [15], [16], [17], [18], [35], [38], [74] enhance generalization ability by exploring elaborate augmentations on pristine videos, with the aim of covering most of the known forgery types. The limitation of these methods is that the augmentation diversity is restricted to known forgeries. Hence, these methods can hardly handle unknown forgeries. Another vein of methods [10], [19], [20], [21], [22], [39], [40] utilize frequency features to improve generalization ability. However, frequency features can easily be disrupted by post-processing such as compression [41]. Inspired by transfer learning, the methods [23], [24], [25], [37] employ zero-shot and few-shot learning to detect new forgeries. Since zero-shot learning cannot access the samples of new forgeries, its performance gain is severely limited. The few-shot learning needs a small portion of samples and corresponding labels of new forgeries. However, although the video-level label is easily obtained, the face-level label is extremely difficult to obtain in practice.

### B. Unsupervised Domain Adaptation

Unsupervised domain adaptation (UDA) aims to address the challenge of transferring knowledge from a source domain

to a target domain when labeled data is scarce or completely absent in the target domain. Ben-David et al. [42] theoretically revealed that the cross-domain common features serve as latent representations that encapsulate shared and domain-common features across diverse domains. The primary objective is to diminish or eliminate domain-specific variations while retaining domain-agnostic information. The acquisition of cross-domain common representation enhances the model's reliability to domain shifts by prioritizing task-relevant information that transcends domain-specific discrepancies. Consequently, the model achieves improved generalization to unlabeled target domains, even in the scenarios with limited available data.

The existing works for addressing UDA can be classified to two main forms, namely, the discrepancy-based approach and the adversarial approach. Concretely, discrepancy-based methods encourage the model to align the domain discrepancy by minimizing the metrics that can measure the distribution discrepancy between the source and target domains [43], [44], [45], [46]. Inspired by the success of generative adversarial network (GAN) [47], recently developed works employed extra adversarial discriminator to align the domain discrepancy, as the feature distributions of source and target domains can be matched by means of confusing the discriminator [48], [49], [50]. In addition, some state-of-the-art methods build up the feature extractor based on modern transformer structure [51], [52], [53], which demonstrates that UDA not only helps traditional CNNs to improve the generalization but also is profitable for transformer-based networks. This motivates us to treat the transformer networks as the cornerstone structure and further explore effective UDA methods for face forgery detection.

Note that the general UDA task targets transferring the knowledge of the semantic class category. By contrast, our approach differs from the aforementioned UDA methods in that we aim to explore the subtle forgery features in the face category only. We also find that the existing adaptation schemes, which only consider the adaptation from the source domain to the target domain, is unlikely to perform well on our task. In contrast, our proposed bi-directional adaptation strategy can further explore the knowledge from the unlabeled data in the target domain, as such mutual adaptation coupled with knowledge transfer with self-distillation enables the model to learn common forgery features across known and new forgeries. To the best of our knowledge, Chen and Tan [54] is the first work that attempted to solve Deepfake detection using unsupervised domain adaptation. However, it is a trivial usage of a naive existing solution without improvement, and hence the detection performance is not satisfied. By contrast, our *DomainForensics* adopts a meticulously designed strategy, named bi-directional adaptation, which can fully learn the common forgery features across domains and it is validated under several practical cross-domain scenarios.

## III. DOMAINFORENSICS

To achieve continuously exposing new forgeries, we formulate DeepFake detection into an unsupervised domain adaptation problem, which transfers the forgery features from



known forgeries to new forgeries, without the need of target labeling. Since the forgery features are very subtle, adapting the general UDA to our task is difficult. As such, we propose a new bi-directional adaptation strategy to fully explore the common forgery features across domains. It is worth noting that our *DomainForensics* is plug and play, i.e., it can be applied to other DeepFake detection architectures.

We start with the problem formulation in Subsection III-A, and then discuss the advantages of our *DomainForensics* over other architectures in Subsection III-B, followed by a performance comparison with existing UDA schemes in Subsection III-C. This naturally motivates us to introduce the new bi-directional adaptation strategy in Subsection III-D and our network architecture in Subsection III-E.

#### A. Problem Formulation

Let the sets of known forgeries and new forgeries corresponding to two different domains be denoted as the source domain  $\mathcal{D}_s$  and target domain  $\mathcal{D}_t$ , respectively. The source domain  $\mathcal{D}_s$  is fully annotated as  $\mathcal{D}_s = \{(x_i^s, y_i^s)\}_{i=1}^{n_s}$ , where  $(x_i^s, y_i^s)$  is the  $i$ -th pair of sample and its corresponding label (i.e., real or not), and  $n_s$  is the number of samples. Differently, the target domain  $\mathcal{D}_t$  only contains samples without any annotations, which are given by  $\mathcal{D}_t = \{x_i^t\}_{i=1}^{n_t}$ , where  $n_t$  is the number of samples. Each domain can be divided into a training set and a testing set as  $\{\mathcal{D}'_s, \mathcal{D}''_s\}$  and  $\{\mathcal{D}'_t, \mathcal{D}''_t\}$ , respectively. We employ  $\mathcal{D}'_s$  and  $\mathcal{D}'_t$  in the training phase and perform evaluation on  $\mathcal{D}''_s$  and  $\mathcal{D}''_t$  in the testing phase. Note that  $\mathcal{D}''_s$  and  $\mathcal{D}''_t$  are unseen during training.

Denote a DeepFake detector as  $\mathcal{H} = \mathcal{G} \circ \mathcal{F}$ , where  $\mathcal{G}$  is the classifier and  $\mathcal{F}$  is the feature extractor. Given a face image  $x$ , the output logits of DeepFake detector  $\mathcal{H}$  can be defined as  $\mathcal{G}(\mathcal{F}(x; \theta_{\mathcal{F}}); \theta_{\mathcal{G}})$ , where  $\theta_{\mathcal{F}}$  and  $\theta_{\mathcal{G}}$  are the parameters of the feature extractor and classifier, respectively. The challenge under this scenario lies in transferring the forgery knowledge learned from known forgeries  $\mathcal{D}_s$  to new forgeries  $\mathcal{D}_t$ , in terms of the underlying marginal distribution discrepancy, i.e., different manipulated approaches. Our goal is to push the feature extractor to learn the common forgery features across different domains, without the supervision of target labels, i.e., achieving favorable performance on both  $\mathcal{D}''_s$  and  $\mathcal{D}''_t$ .

#### B. DomainForensics Versus Existing Architectures

Using data augmentation is the most typical solution to improve the generalizability of detection methods, e.g., FWA [55], Face X-ray [16], SBI [18]. These methods attempt to synthesize various pseudo-fake faces to cover known forgery as much as possible. By training on these augmented samples, the detectors can learn the common forgery features. In this scenario, the manipulation operations of new forgeries should be known as a prior, in order to synthesize the applicable pseudo-fake faces. However, the technical details of new forgeries may not be accessible in reality, limiting the application of these existing methods. In our scenario, we do not require the technical details of new forgeries. Instead, we can collect a set of videos that contains new forgery faces, and simply extract all faces in a video without knowing the

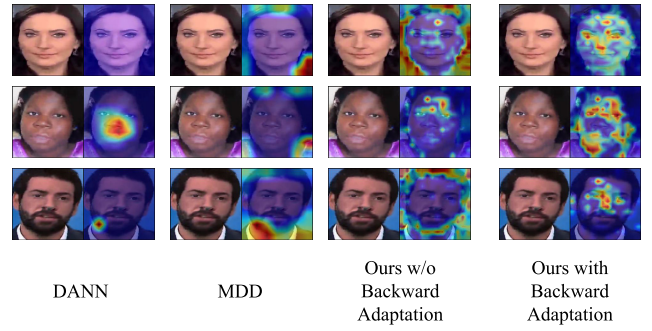


Fig. 2. Grad-CAM visualization. We train the models, including DANN [48], MDD [46] and our *DomainForensics*, and visualize the activation maps on FF++ dataset under FS→F2F scenario. These figures show that models fails to fully capture the common forgery features when only employing one-directional adaptation.

label (real or fake) of faces. Using these samples, we can align the detectors to learn the transferable knowledge from the known forgery to this new forgery.

This scenario is practical and useful. For example, we can obtain video sets by searching the keywords, e.g., DeepFake, on video platforms. The obtained videos are likely a mix of real and fake faces due to the natural deviation of search engines, i.e., we cannot ensure whether a video that appeared in search results is real or fake. Even though the search results are perfectly matched, i.e., the video-level annotation is correct, the obtained faces can still be mixed, as a video usually contains multiple faces and we cannot know which face is real or fake if only video-level annotation is given. Under this practical circumstance, our method can expose new forgeries with these unlabeled samples.

#### C. Comparison With Existing UDA Schemes

To learn domain-common forgery features, we consider building up a solution from the perspective of UDA. With the aim of reducing the domain discrepancy, early domain adaptation methods put eyes on transferring knowledge from the source domain to the target domain [45], [46], [48]. However, such one-direction adaptation methods are insufficient in digging out the subtle forgery knowledge from unlabeled target data. Fig. 2 shows several examples of the feature activation maps visualized by Grad-CAM [56] using one-direction adaptation. The first two columns from left to right show the CAMs of DANN [48] and MDD [46], two typical domain adaptation methods. The last two columns are the CAMs of our approach without and with backward adaptation. It can be seen that the typical one-direction domain adaptation cannot locate the forgery features very well in these examples, either paying attention to the background area or the central local face part. In comparison, the CAMs of our method trained with backward adaptation scatter all around the whole face, activating more correct forgery features than these methods including ours without using backward adaptation. This illustration demonstrates that using existing domain adaptation methods for our task is questionable and inspires us to develop a devoted solution, which we detail in the following subsection.

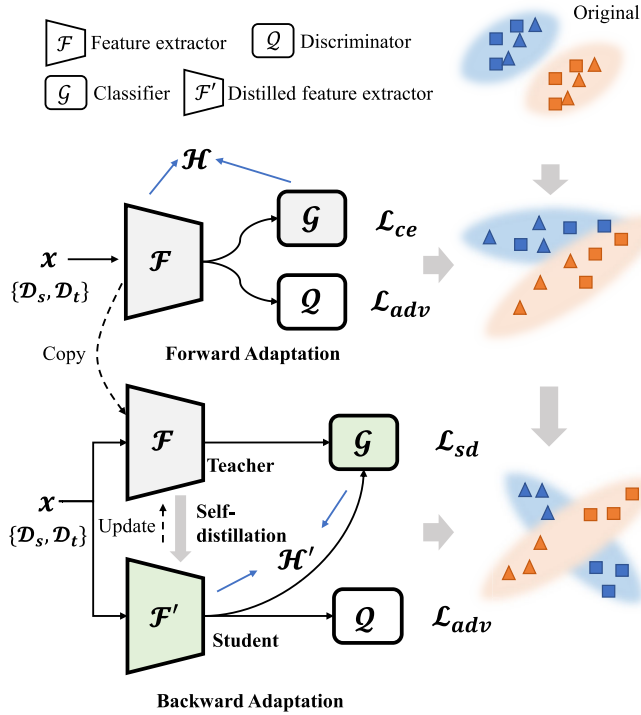


Fig. 3. Illustration of the proposed bi-directional adaptation strategy, containing forward adaptation and backward adaptation with  $\mathcal{H}'$  as the final DeepFake detector. Note that other architectures can also be used in our framework.

#### D. Bi-Directional Adaptation

As aforementioned, to learn domain-common forgery features, we consider building up a solution from the perspective of UDA. Existing general domain adaptation usually considers one-direction to transfer knowledge from the source domain to the target domain. However, such an one-direction adaptation is insufficient in learning transferable knowledge, as it neglects to learn from unlabeled target data, as demonstrated in the previous subsection. Thus, we propose a new bi-directional adaptation strategy, which consists of both forward adaptation and backward adaptation. The forward adaptation stage aims to transfer the knowledge from the source domain to the target domain, just as the existing solutions [45], [48], [49]. However, due to the limitation of such an one-direction adaptation, a portion of target domain knowledge is lost in the transfer, thus hindering the performance of the target domain. To eliminate or mitigate this deficiency, we develop a backward adaptation stage to fine-tune the DeepFake detector on the target domain, while retaining the learned knowledge in the forward adaptation stage. Fig. 3 illustrates the proposed bi-directional adaptation strategy.

1) *Forward Adaptation*: In this stage, we aim to learn the common forgery features by adapting the source domain to the target domain. Concretely, we design two loss terms. The first term is a cross-entropy loss  $\mathcal{L}_{ce}$ , which trains the DeepFake detector on the fully-annotated source domain:

$$\min_{\theta_{\mathcal{F}}, \theta_{\mathcal{G}}} \mathcal{L}_{ce} = \mathbb{E}_{(x_i^s, y_i^s) \sim \mathcal{D}_s} [-\log \mathcal{G}(\mathcal{F}(x_i^s; \theta_{\mathcal{F}}); \theta_{\mathcal{G}})^{y_i^s}]. \quad (1)$$

This loss term enables the DeepFake detector to distinguish classes, i.e., telling apart real and fake faces. To transfer this

knowledge from the source domain to the target domain, the essence of this stage is to push the feature extractor  $\mathcal{F}$  to generate features that cannot be identified from which domain they come. To achieve this goal, we design an adversarial loss  $\mathcal{L}_{adv}$  as the second loss term, which guides the training of feature extractor  $\mathcal{F}$  with a discriminator  $\mathcal{Q}$  in an adversarial manner. Denote  $\theta_{\mathcal{Q}}$  as the parameters of discriminator  $\mathcal{Q}$ . This loss term can be defined as

$$\min_{\theta_{\mathcal{F}}} \max_{\theta_{\mathcal{Q}}} \mathcal{L}_{adv} = \mathbb{E}_{x_i^s \sim \mathcal{D}_s} [\log \mathcal{Q}(\mathcal{F}(x_i^s; \theta_{\mathcal{F}}); \theta_{\mathcal{Q}})] + \mathbb{E}_{x_i^t \sim \mathcal{D}_t} [\log (1 - \mathcal{Q}(\mathcal{F}(x_i^t; \theta_{\mathcal{F}}); \theta_{\mathcal{Q}}))]. \quad (2)$$

The discriminator  $\mathcal{Q}$  outputs binary labels, i.e.,  $\mathcal{Q}(\cdot; \theta_{\mathcal{Q}}) \in \{0, 1\}$ , where 0 denotes target domain and 1 denotes source domain. The overall loss of this stage is written as

$$\mathcal{L}_{fas} = \alpha_1 \cdot \mathcal{L}_{ce} + \alpha_2 \cdot \mathcal{L}_{adv}, \quad (3)$$

where  $\alpha_1$  and  $\alpha_2$  are the weight factors. The training of this stage is an adversarial min-max game between feature extractor  $\mathcal{F}$  and discriminator  $\mathcal{Q}$ . To update  $\mathcal{F}$ , we fix the parameters of  $\mathcal{Q}$ , and vice versa. Note that in the optimization of  $\mathcal{L}_{adv}$ , we do not need any class labels from both domains.

2) *Backward Adaptation*: The existing adaptation methods usually only consider the one-direction forward adaptation and their performance are limited on our task due to the loss of knowledge in transferring. The key to overcome this limitation is to mine effective forgery features more specific to the target domain. Concretely, we reverse the adaptation by training the DeepFake detector on the target domain and then transferring the knowledge from the target domain back to the source domain. The major challenge here is that no labels are provided in the target domain, and thus we cannot refine the DeepFake detector in the form of cross-entropy loss  $\mathcal{L}_{ce}$  that is used in the forward adaptation stage.

To address this difficulty, we adopt self-distillation in our framework to further explore the specific representations of the target domain. Inspired by SIMCLRv2 [27], we employ a teacher-student network model as the training structure. Specifically, during the training process, we adopt the feature extractor  $\mathcal{F}$  with the parameters  $\theta_{\mathcal{F}}$  and the classifier  $\mathcal{G}$  with the parameters  $\theta_{\mathcal{G}}$  used in the forward adaptation stage as the teacher model. We then create a new feature extractor  $\mathcal{F}'$  with the parameters  $\theta_{\mathcal{F}'}$ , which has the same model structure as  $\mathcal{F}$  and whose parameters  $\theta_{\mathcal{F}'}$  are initialized with the same parameters as  $\mathcal{F}$ .  $\mathcal{F}'$  is combined with the classifier  $\mathcal{G}$  to make up the student model. It can be seen that the teacher and student models use the same model structure.

Concretely, the self-distillation loss  $\mathcal{L}_{sd}$  is defined as

$$\min_{\theta_{\mathcal{F}}, \theta_{\mathcal{F}'}, \theta_{\mathcal{G}}} \mathcal{L}_{sd} = -\sum_{x_i^t \in \mathcal{D}_t} \left( \sum_k \mathcal{P}(k|x_i^t; \tau) \log \mathcal{P}'(k|x_i^t; \tau) \right). \quad (4)$$

Here  $\mathcal{P}$  and  $\mathcal{P}'$  are the distillation probabilities of the teacher model and student model, respectively. In particular,

$$\mathcal{P}(k|x_i; \tau) = \frac{\exp(\mathcal{G}(\mathcal{F}(x_i; \theta_{\mathcal{F}}); \theta_{\mathcal{G}})^k / \tau)}{\sum_{k'} \exp(\mathcal{G}(\mathcal{F}(x_i; \theta_{\mathcal{F}}); \theta_{\mathcal{G}})^{k'} / \tau)}, \quad (5)$$

**Algorithm 1** Training Procedure of Bi-Directional Adaptation

**Input:** Source domain  $\mathcal{D}'_s$ ; Target domain  $\mathcal{D}'_t$ ; Initial feature extractor  $\mathcal{F}$  with parameters  $\theta_{\mathcal{F}}$ ; Initial classifier  $\mathcal{G}$  with parameters  $\theta_{\mathcal{G}}$ ; Initial distilled feature extractor  $\mathcal{F}'$  with parameters  $\theta_{\mathcal{F}'}$ ; Initial discriminator  $\mathcal{Q}$  with parameters  $\theta_{\mathcal{Q}}$ ; Number of forward adaptation epochs  $T_1$ ; Number of backward adaptation epochs  $T_2$

// Forward adaptation stage

**while**  $t \leq T_1$  **do**  $\triangleright t = 0$

  // Fix  $\mathcal{Q}$

$\min_{\theta_{\mathcal{F}}, \theta_{\mathcal{G}}} \alpha_1 \cdot \mathcal{L}_{ce} + \alpha_1 \cdot \mathcal{L}_{adv}$

  // Fix  $\mathcal{F}, \mathcal{G}$

$\max_{\theta_{\mathcal{Q}}} \mathcal{L}_{adv}$

**end while**

// Backward adaptation stage

**while**  $t \leq T_2$  **do**  $\triangleright t = 0$

  // Fix  $\mathcal{Q}$

$\min_{\theta_{\mathcal{F}}, \theta_{\mathcal{F}'}, \theta_{\mathcal{G}}} \alpha_3 \cdot \mathcal{L}_{sd} + \alpha_4 \cdot \mathcal{L}_{adv}$

  // Fix  $\mathcal{F}, \mathcal{F}', \mathcal{G}$

$\max_{\theta_{\mathcal{Q}}} \mathcal{L}_{adv}$

  // Update  $\mathcal{F}$

$\theta_{\mathcal{F}} = \theta_{\mathcal{F}'}$

**end while**

**Output:**  $\mathcal{F}', \mathcal{G}$

where  $k$  is the class index and  $\tau$  is a scalar temperature parameter.  $\mathcal{P}'(k|x_i; \tau)$  is obtained by replacing  $\mathcal{F}$  with  $\mathcal{F}'$  in (5). The Eq.(4) is designed to distill the knowledge from the teacher model learned in the target domain into the student model. By using self-distillation, the student model can enhance the features of the teacher model, thus capturing more effective knowledge regarding the target domain in comparison to the teacher model. We then utilize  $\mathcal{L}_{adv}$  as described in the forward adaptation stage on the same discriminator  $\mathcal{Q}$  and the distilled feature extractor  $\mathcal{F}'$ . The overall loss of this stage is the combination of these two loss terms as

$$\mathcal{L}_{bas} = \alpha_3 \cdot \mathcal{L}_{sd} + \alpha_4 \cdot \mathcal{L}_{adv}, \quad (6)$$

where  $\alpha_3$  and  $\alpha_4$  are weight factors. We optimize this loss adversarially as in the forward adaptation stage. At the end of each training epoch, we update the teacher model by copying the parameters of the student model to the teacher model:

$$\theta_{\mathcal{F}} = \theta_{\mathcal{F}'}. \quad (7)$$

This backward adaptation strategy provides more accurate guidance from the teacher model, thus promoting the student model to learn more knowledge. Both the self-distillation loss  $\mathcal{L}_{sd}$  and adversarial loss  $\mathcal{L}_{adv}$  do not need any class labels from both domains.

3) *Training and Inference:* The training procedure of our framework is summarized in Algorithm 1. In inference, we use feature extractor  $\mathcal{F}'$  and classifier  $\mathcal{G}$  as our DeepFake detector  $\mathcal{H}' = \mathcal{G} \circ \mathcal{F}'$ .

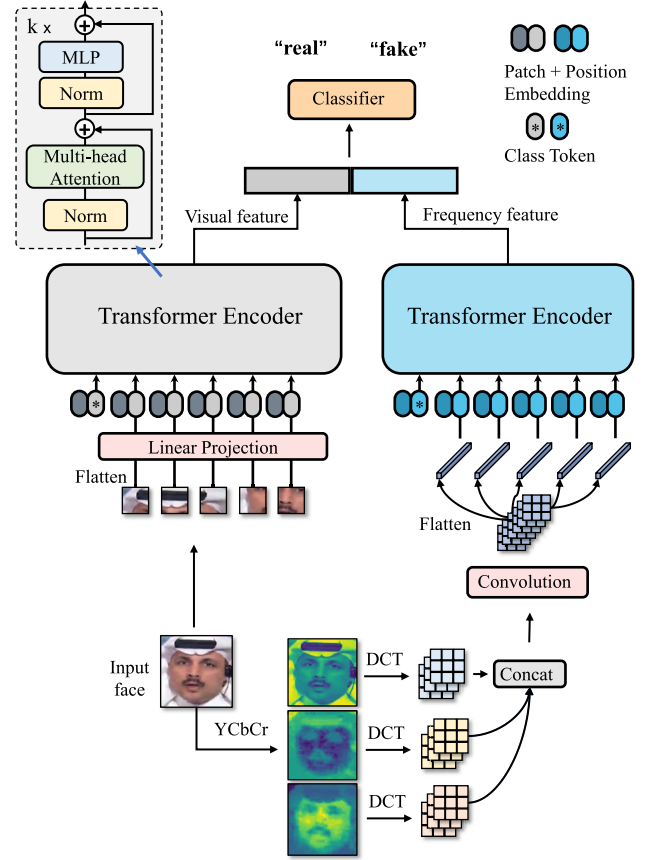


Fig. 4. Network architecture for DeepFake detector.

### E. Network Framework

Fig. 4 depicts the network architecture for our DeepFake detector. We design a ViT-based network as our feature extractor  $\mathcal{F}$  due to its strong power on vision tasks. Specifically, our feature extractor has two branches, a visual branch and a frequency branch. For the visual branch, we split the face image into 196 patches. These patches are flattened and linear transformed to patch embeddings, which are then equipped with position embeddings as the tokens for a ViT [29] architecture to extract visual features. The ViT in this branch contains  $l$  transformer layers, each of which is composed by a multi-headed self-attention (MSA) layer and a MLP layer [57]. For the frequency branch, we first convert the face image from RGB color space to YCbCr color space and then apply DCT transformation to each component with a  $8 \times 8$  block [10]. The transformed frequency maps are concatenated together and sent into a convolution block. We then flatten the feature maps from the convolution block into 196 vectors, which are used as the input to another ViT architecture for frequency feature extraction. This ViT architecture contains  $m$  transformer layers. The visual features and frequency features are concatenated together as the forgery features for DeepFake detection. The classifier  $\mathcal{G}$  has a simple structure with only two linear layers, which takes the forgery features as input and outputs the logit of prediction. It is worth emphasizing that our method is independent of the network architecture, and can be integrated into other mainstream architectures. We use



the architecture of Fig. 4 in our experiments as it can achieve the best performance, which will be confirmed by Table IX in Ablation Study.

#### IV. EXPERIMENTS

##### A. Experimental Settings

1) *Datasets*: We evaluate our approach on three public deepfake detection datasets, which are FF++ [13], Celeb-DF [14], StyleGAN [58], DFDCP, and FFIW.

FF++ is a widely used dataset in deepfake detection. It includes 1,000 original videos from YouTube, covering a wide range of subjects and scenarios, and consists of four different manipulation techniques: Deepfakes (DF), Face2Face (F2F), FaceSwap (FS), and NeuralTextures (NT), each representing a distinct form of facial manipulation. All these videos have three compression versions, raw, high quality (HQ) and low quality (LQ). From these original videos, 720 videos, 140 videos and 140 videos are used for training, validating and testing, respectively. For each manipulation technique, we employ the same partition as the original videos for training, validating, and testing. We focus on the frame-level deepfake detection and perform evaluation on both HQ and LQ data. As our approach does not rely on extra augmentation operations, we only crop faces from each frame during the preprocessing stage. Concretely, we randomly extract 8 frames for each video clip and crop out the face with  $1.3\times$  of the detection box obtained by RetinaFace [59].

Celeb-DF was proposed more recently, which provides a diverse set of challenges, such as pose, lighting conditions, facial expressions and camera quality, commonly encountered in real-world scenarios. It also offers different levels of manipulation difficulty, ranging from subtle and realistic manipulations to more obvious and noticeable ones. This dataset contains 590 original videos and 5,639 DeepFake videos. We respectively use 5,710 and 518 videos for training and testing. To construct the training data, we randomly extract 32 frames from real videos and 4 frames from fake videos for data balance. As for the testing set, we extract 16 frames from real and fake videos. The face is cropped out using the same way as for FF++.

StyleGAN [58] trains a generative adversarial network on Flickr-Faces-HQ dataset (FFHQ), which consists of 70,000 real face images, to synthesize high-quality fake faces. The synthesizing process focuses on human faces and offers various high-quality synthetic faces by disentangling style and content information in latent space to control the process of face generation. We randomly select 2,500 original images from FFHQ and 2,500 synthetic images from the generated dataset, where 2,000 original images and 2,000 synthetic images are for training and the remaining images are for testing.

DFDCP [71] is a preview dataset for deepfakes detection challenge consisting of 5,000 videos with two facial modification algorithms, considering diversity in several axes (gender, skin-tone, age, etc.).

FFIW-10K [67] is a large-scale dataset, which comprises 10,000 pristine videos and 10,000 high-quality forgery videos

for training. Another 250 real videos and 250 manipulated videos are also provided for testing. Different from the aforementioned datasets, this dataset contains multiple forgery faces in a single video.

2) *Domain Adaptation Scenarios*: Since a new forgery can either be generated by a new manipulation method or come from an unseen data distribution, we design three adaptation scenarios, which are *Cross-manipulation-methods*, *Cross-manipulation-datasets* and *Cross-manipulation-types*, respectively. Cross-manipulation-methods is the case of crossing different manipulation methods in the same type (e.g., faceswap). Cross-manipulation-datasets represents the adaptation from one dataset to another different one. Cross-manipulation-types is to adapt one type of forgery to a different type. This scenario is more challenging, as different types have significant discrepancies in the manipulation methods, datasets and forged areas, e.g., faceswap faces to GAN-synthesized faces.

To make a fair comparison with the previous works [16], [35], each domain adaptation scenario is trained on the training set and tested on the testing set. Take the Cross-manipulation-methods scenario on FF++ dataset as an example. In training, we use the training set of source and target manipulation methods. In inference, we evaluate the trained model on the testing set of source and target manipulation methods. This configuration also applies to the other two scenarios.

3) *Implementation Details*: Following the design in Vit-Base [29], the feature extractor contains 12 (i.e.,  $l = 12$ ) transformer blocks, while we use 4 (i.e.,  $m = 12$ ) transformer blocks in frequency branch. The input images are resized to (224, 224). The weights of loss terms are set to  $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4 = 1$ . The training batch size is set to 32 in forward adaptation and 24 in backward adaptation, as the teacher model occupy extra GPU memory. In the forward adaptation stage, we employ SGD as the optimizer, where the learning rate, momentum, and weight decay are set to 0.001, 0.9, and 0.0005, respectively. We stop the training of forward adaptation after 20 epochs. For backward adaptation, we set the learning rate to 0.0001 and the number of training epochs to 10. The temperature is set to  $\tau = 0.5$ . All experiments are conducted using Pytorch [60] on two Nvidia GTX 2080Ti GPUs.

##### B. Results of Cross-Manipulation-Methods

We investigate two settings in this scenario: one-to-one adaptation (O2O) and one-to-many adaptation (O2M). O2O uses one method as source and another one as target, while O2M uses one method as source and many other methods as target, which is more practical, as the collected target videos may have many new forgeries. All the experiments in this part are conducted on FF++ dataset and evaluated using the area under curve (AUC) metric.

1) *Performance of O2O and O2M*: In O2O, we select one of the four manipulation methods (DF, F2F, NT, FS) as the source domain and select a different method as the target domain. The top part of Table I shows the performance of our method under this O2O setting on the HQ set. Gray and yellow coloured values denote the performance on the source domain and

TABLE I  
THE PERFORMANCE OF OUR PROPOSED METHOD UNDER TWO  
CROSS-MANIPULATION SCENARIOS ON FF++. THE TOP  
PART IS THE RESULTS OF O2O AND THE BOTTOM  
PART IS THE RESULTS OF O2M

Adaptation	DF	F2F	NT	FS
Baseline	100.0	69.43	69.55	55.18
DF → F2F	99.00	99.08	66.21	64.67
DF → NT	87.88	60.02	95.03	43.49
DF → FS	98.82	75.68	57.41	99.37
Baseline	81.71	99.13	55.84	69.74
F2F → DF	99.77	98.67	73.56	63.55
F2F → NT	89.89	95.59	95.85	57.11
F2F → FS	90.49	98.86	51.16	99.89
Baseline	79.74	55.19	95.72	44.98
NT → DF	99.28	76.79	95.80	42.94
NT → F2F	92.79	97.44	93.27	63.74
NT → FS	78.96	63.28	70.05	97.81
Baseline	84.34	71.06	43.63	99.53
FS → DF	98.90	79.62	61.16	98.95
FS → F2F	85.40	99.14	50.67	99.69
FS → NT	88.74	73.62	96.27	95.71
DF → F2F,NT,FS	97.64	98.64	88.73	99.12
F2F → DF,NT,FS	99.49	97.38	93.02	98.15
NT → DF,F2F,FS	98.90	97.20	94.52	95.35
FS → DF,F2F,NT	99.26	99.18	86.91	99.00

target domain, respectively. “Baseline” denotes training the proposed architecture on source domain and directly testing on other domains without bi-directional adaptation. The results reveal that our method performs favourably on both source and target domains, achieving the AUC metric 95%+ on most of the source and target adaptation pairs. This demonstrates the effectiveness of our method. We also observe that the performance with NT → FS is limited on the source. This is probably due to the large gap between NT and FS, and thus part of the knowledge in the source domain is lost.

In O2M, we select one manipulation method as source domain and use the other three methods as target domain. The performance achieved by our method are shown in the bottom part of Table I. We can observe that our method performs well on all the manipulation methods, which demonstrates that our method can learn the common forgery features even if the target domain is mixed with different manipulations.

Compared to the O2O scenario, O2M is more practical, as the daily emerged videos likely contain various manipulation methods. However, the foundation of O2M is O2O, as it still attempts to find the common knowledge among these various manipulation methods. Thus the O2O setting is also important, as it serves as the basis for the improvement of O2M.

2) *Comparison With State-of-the-Arts*: We compare our method with several state-of-the-art methods, including the augmentation-based methods (Face X-ray [16], SBI [18]), frequency-based methods (SRM [21], FDFL [10]), transfer-based methods (FOT [23], DDT [24], FEAT [54]), and other methods (Xception [13], MATT [61], LTW [34], RECCE [62], SOLA-sup [63], FTCN [64], TALL-Swin [73]). Note that TALL-Swin has no reports corresponding to this scenario.

Thus we reproduce its results following its default settings. This method is trained using all videos in FF++.

We would like to clarify that the experiment configurations between our method and these other methods are different, as these methods tackle this task under different scenarios. Specifically, these data-level methods first obtain and empirically analyze the new forgery videos, and then summarize a common forgery knowledge as a prior, e.g., the blending boundary and frequency clues. By contrast, the prior knowledge for us is the collected videos without labels. Even though the experiment configurations are not the same, the results can also reflect the effectiveness of our method in exposing new forgeries.

Table II compares the performance of our method with those of the existing methods under the O2O scenario on both HQ and LQ levels, where the results of our method are marked by gray colour. For the existing methods, we use the score of each method reported in its original paper. As SBI [18] was only evaluated on raw quality images, we retrain it using the code provided in [18] on HQ and LQ for a fair comparison.

For the HQ level, it can be seen from the top part of Table II that our method outperforms the augmentation-based methods, Face X-ray and SBI, by a large margin. Since these two methods require prior knowledge of manipulation, they can hardly handle the forgery that has notable differences from the prior knowledge. The frequency-based method SRM performs well on several cases, notably, FS → F2F and NT → F2F. But its performances are highly degraded in many more cases compared to our method. For example, in DF → FS, our method is nearly 100% better than SRM, and in FS → DF, our method is 44% better than SRM, while in DF → F2F, our method is 30% better than SRM, and in F2F → DF, our method is 20% better than SRM. Similar observations can be drawn by comparing our method with SOLA-sup. Based on the available data from [64], the performance of our method are slightly better than those of FTCN. When compared to TALL-Swin, we observe our method achieves better generalization performances than TALL-Swin. It can also be seen that our method significantly outperforms LTW, Xception, FOT, DDT and FEAT. Moreover, in the more challenging LQ level, our method significantly outperforms all the five benchmarks, as can be seen from the bottom part of Table II.

Fig. 5 shows the t-SNE [65] visualizations of feature distribution for two O2O cross-manipulations (FS → F2F, FS → NT) without and with adaptation using our method on both source and target domain. For the target domain, it can be seen that the features of real and fake faces separate well using our method, in contrast to the mixed features without our method (see (a,b) and (c,d)). For the source domain, we can observe that their feature distributions are still discriminative (see (e,f) and (g,h)). Fig. 6 shows several examples of Grad-CAM [56] (the left six columns), which reveals that our method catches more discriminative features on face regions for both real and fake faces.

3) *Comparison With Existing UDA Methods*: As our approach draws inspiration from UDA to enhance the generalization ability when facing unseen DeepFake techniques, we also perform a further comparison with the recently



TABLE II

COMPARISON OF OUR METHOD WITH STATE-OF-THE-ARTS UNDER O2O CROSS-MANIPULATION SCENARIO AT HQ LEVEL (TOP) AND LQ LEVEL (BOTTOM) IN FF++. THE BOLD NUMBER DENOTES THE BEST PERFORMANCE AND THE UNDERLINED ONE DENOTES THE SECOND BEST

Method	DF			F2F			NT			FS		
	F2F	NT	FS	DF	NT	FS	F2F	DF	FS	F2F	NT	DF
(HQ) Face X-ray [16]	63.30	69.80	60.00	63.00	94.50	93.80	91.70	70.50	91.00	96.10	95.70	45.80
(HQ) SBI [18]	67.08	62.14	72.82	91.63	62.14	72.82	67.08	91.63	72.82	67.08	62.14	91.63
(HQ) LTW [34]	80.02	77.30	64.00	92.70	77.30	64.00	80.02	92.70	64.00	80.02	77.30	92.70
(HQ) Xception [13]	73.60	73.60	49.00	80.30	69.60	76.20	81.30	80.00	73.10	88.80	71.30	66.40
(HQ) SRM [21]	76.40	81.40	49.70	83.70	<b>98.40</b>	<b>98.70</b>	<b>99.50</b>	89.40	<b>99.30</b>	<b>99.30</b>	<b>98.00</b>	68.50
(HQ) SOLA-sup [63]	<u>97.29</u>	<b>98.48</b>	69.72	99.73	96.02	93.50	<u>97.69</u>	<b>99.64</b>	<b>99.76</b>	98.13	92.07	<b>99.11</b>
(HQ) FTCN [64]	-	-	-	98.00	96.00	95.90	-	-	-	-	-	-
(HQ) FOT [23]	-	-	-	-	-	72.57	-	-	-	-	-	-
(HQ) DDT [24]	-	64.10	-	-	-	-	-	78.82	-	-	-	-
(HQ) FEAT [54]	-	-	<u>88.62</u>	-	-	-	-	-	-	-	-	-
(HQ) TALL-Swin [73]	52.96	52.11	77.36	62.55	60.49	55.54	64.92	62.09	52.20	50.55	51.96	85.08
(HQ) Baseline	69.43	69.55	55.18	81.71	55.84	69.74	55.19	79.74	44.98	71.06	43.63	84.34
(HQ) <b>Baseline+BA (Ours)</b>	<b>99.08</b>	95.03	<b>99.37</b>	<b>99.77</b>	95.85	<b>99.89</b>	97.44	99.28	97.81	99.14	96.27	98.90
(LQ) SBI [18]	70.55	65.95	68.25	85.60	65.95	68.25	70.55	85.60	68.25	70.55	65.95	85.60
(LQ) FDFL [10]	58.90	63.61	66.87	67.55	55.35	66.66	74.21	79.09	74.21	54.64	49.72	75.90
(LQ) LTW [34]	<u>72.40</u>	60.80	68.10	75.60	60.80	68.10	72.40	75.60	68.10	<u>72.40</u>	60.80	75.60
(LQ) MATT [61]	66.41	66.01	67.33	73.04	71.88	65.10	80.61	74.56	60.90	61.65	54.79	82.33
(LQ) RECCE [62]	70.66	67.34	74.29	75.99	72.32	64.53	80.89	78.83	63.07	64.44	56.70	82.39
(LQ) TALL-Swin [73]	55.50	52.93	74.39	63.64	57.82	53.15	68.65	63.19	51.42	50.44	53.11	83.91
(LQ) Baseline	61.18	59.49	71.17	65.88	51.43	59.07	56.36	65.33	52.49	56.25	44.38	70.45
(LQ) <b>Baseline+BA (Ours)</b>	<b>92.13</b>	<b>86.89</b>	<b>98.09</b>	<b>93.52</b>	<b>79.64</b>	<b>94.22</b>	<b>86.96</b>	<b>96.38</b>	<b>93.43</b>	<b>88.16</b>	<b>80.85</b>	<b>99.02</b>

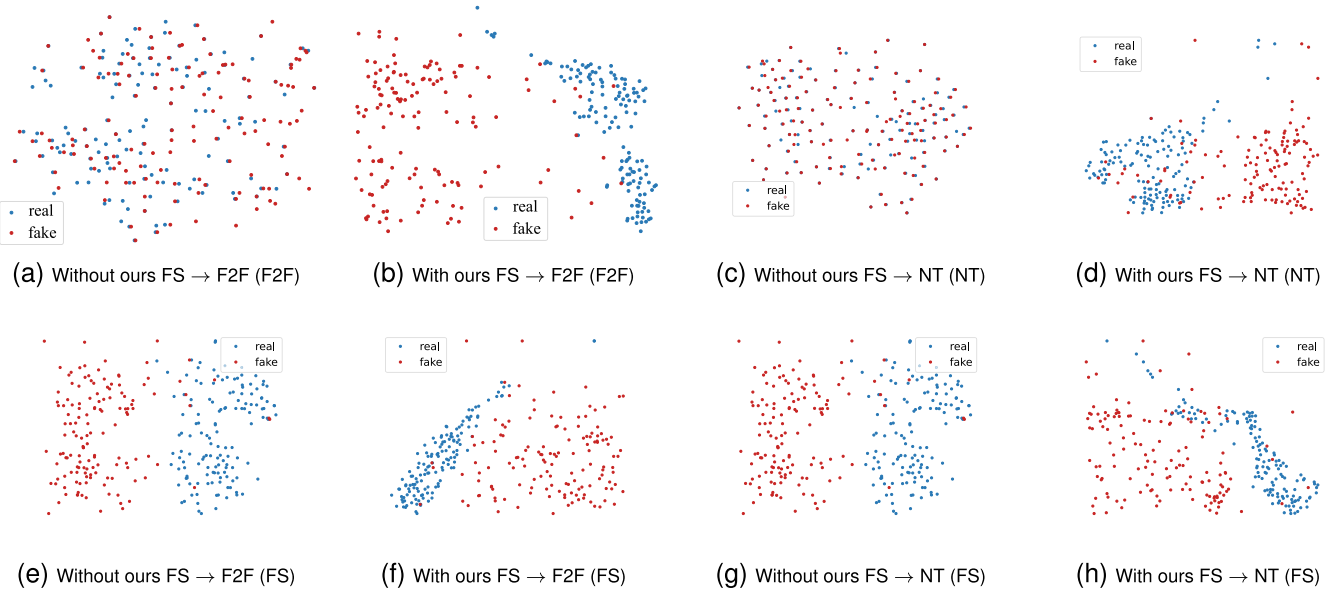


Fig. 5. T-SNE [65] visualization on FF++ (HQ).

developed domain adaptation methods, SSRT [53], MDD [46] and DANN [48]. Specifically, we retrain their models on FF++ HQ and LQ scenarios by using their published codes. The experimental results are summarized in Table III, which demonstrate that our method outperforms these three domain adaptation methods in recognizing faces with unseen forgeries. SSRT [53] is a counterpart method that also employs a powerful vision transformer as the backbone network. It can be seen that SSRT occasionally fails to transfer knowledge in some scenarios, e.g., DF→NT and FS→F2F, while our bi-directional network achieves more stable and superior

performance than SSRT almost in all adaptation tasks. DANN and MDD respectively represent an adversarial-based method and a discrepancy-based method. We observe that DANN achieves more generalized performance than MDD but our approach surpasses both DANN and MDD by a large margin. The reason mainly lies in that the proposed bi-directional adaptation strategy enables the model to fully grab the common forgery features across domains. Overall, the results suggest that our approach effectively encourages the detector to learn common forgery features, leading to improved generalization ability.

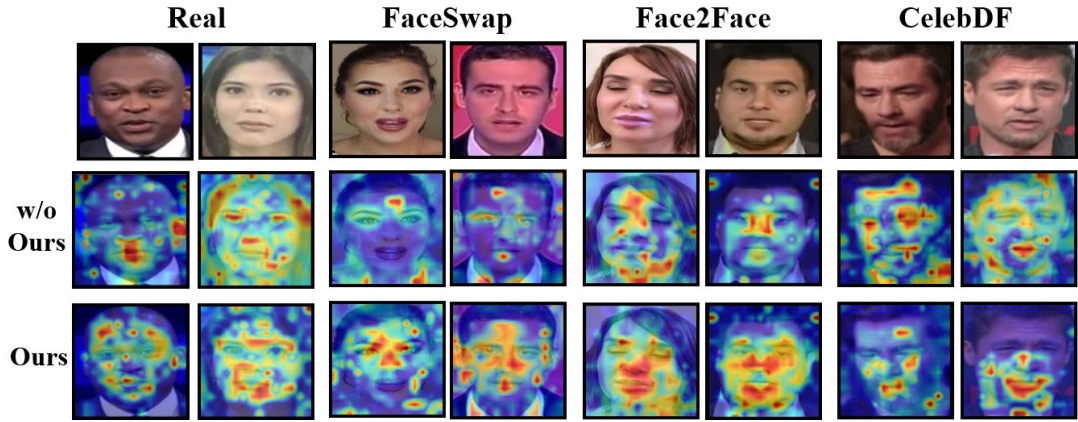


Fig. 6. Grad-CAM [56] visualization without and with adaptation using our method.

TABLE III

COMPARISON OF OUR METHOD WITH STATE-OF-THE-ARTS DOMAIN ADAPTATION METHODS. THE BOLD NUMBER DENOTES THE BEST PERFORMANCE

Method	DF			F2F			NT			FS		
	F2F	NT	FS	DF	NT	FS	F2F	DF	FS	F2F	NT	DF
(HQ) SSRT [53]	51.43	50.36	61.79	86.79	49.64	50.36	54.64	90.71	49.29	58.21	50.00	86.43
(HQ) MDD [46]	71.79	73.57	61.79	80.36	70.00	80.71	77.50	88.93	75.00	75.71	50.37	73.93
(HQ) DANN [48]	95.71	90.71	94.29	94.99	88.57	96.43	90.36	92.86	89.64	91.43	87.14	95.00
(HQ) <b>Ours</b>	<b>99.08</b>	<b>95.03</b>	<b>99.37</b>	<b>99.77</b>	<b>95.85</b>	<b>99.89</b>	<b>97.44</b>	<b>99.28</b>	<b>97.81</b>	<b>99.14</b>	<b>96.27</b>	<b>98.90</b>
(LQ) SSRT [53]	50.35	50.40	56.79	73.21	54.28	56.43	66.79	68.93	59.29	50.36	50.14	92.14
(LQ) MDD [46]	62.50	57.14	62.14	68.93	56.79	62.50	70.71	77.49	65.71	59.29	49.99	77.50
(LQ) DANN [48]	78.21	71.43	83.93	82.86	66.78	81.79	77.50	81.07	75.71	80.36	68.21	90.36
(LQ) <b>Ours</b>	<b>92.13</b>	<b>86.89</b>	<b>98.09</b>	<b>93.52</b>	<b>79.64</b>	<b>94.22</b>	<b>86.96</b>	<b>96.38</b>	<b>93.43</b>	<b>88.16</b>	<b>80.85</b>	<b>99.02</b>

### C. Results of Cross-Manipulation-Datasets

In this experiment, we use FF++ as source domain and Celeb-DF as target domain. Table IV compares our method with the state-of-the-art methods, F<sup>3</sup>-Net [19], MATT [61], RECCE [62], Two-branch [66], LTW [34], DSP-FWA [15], MTD-Net [36], Xception [13], CFFs [37], DAM [67], SOLA-sup [63], SPSL [20], LiSiam [38], SLADD [35], Lip-Forensics [68], HFI-Net [39], FTCN [64], AltFreezing [72], TALL [73], F<sup>2</sup>Trans [40], and PCL+I2G [17]. It can be seen from Table IV that our method achieves 90.22% at HQ and 81.39% at LQ, respectively, outperforming most other methods by a large margin. Note that PCL+I2G [17] synthesizes forged faces with raw quality data since the uncompressed data can provide more distinct features than the compression ones. Our method outperforms TALL-ViT [73], which uses spatio-temporal modeling for learning local and global contextual deepfake patterns, when employing a vision transformer as the backbone network. Additionally, it achieves comparable performance with TALL-Swin. This demonstrates that our method is capable of transferring the forgery knowledge from one dataset to another. The last two columns of Fig. 6 depict the examples of Grad-CAM from FF++ to Celeb-DF, which shows that our method concentrates more on forgery regions than the case without using our method.

For a comprehensive study, we then conduct a more challenging scenario, which uses the FF++ as the source domain and the DFDCP [71] or FFIW [67] as the target domain.

The results are presented in Table V. We observe that our method achieves better performance than others on the DFDCP dataset, mainly because our method can learn common forgery features among different manipulated methods. However, SBI surpasses our method on the FFIW dataset. This is because this dataset contains many side faces, while the faces in FF++ dataset are usually frontal. This discrepancy largely increases the domain gap, which may disturb our model in learning common forgery features. Since SBI is designed to create a variety of faces for data augmentation, which likely includes such distorted faces in training, resulting in better performance than us.

### D. Results of Cross-Manipulation-Types

In this part, we investigate the feasibility of our method to adapt faceswap faces to GAN-synthesized faces. Faceswap replaces the central region of face and retains other regions unchanged, while GAN synthesizes the whole face image. We use each of faceswap methods in FF++ [13] as source and adapt it to the StyleGAN dataset [58]. Table VI shows the performance of our method in comparison to the three benchmark methods, MATT [61], RECCE [62] and SBI [18], under this setting. For fair comparison, we use the codes provided in [18], [61], and [62] for SBI, MATT and RECCE under their default settings. Observe that these three methods perform poorly, and in most cases their performance are below 50%. This is because these methods are designed to detect faceswap

TABLE IV

PERFORMANCE COMPARISON OF OUR METHOD WITH STATE-OF-THE-ARTS UNDER CROSS-DATASETS SCENARIO FROM FF++ TO CELEB-DF. TRAINING SET FF++ INDICATES FF++ WITH FOUR MANIPULATION METHODS (DF,F2F,NT,FS) AS SOURCE, AND FF++(DF) MEANS FF++ WITH ONLY ONE MANIPULATION METHOD DF AS THE TRAINING SET

Method	Training Set	FF++	Testing Set Celeb-DF
(HQ) LTW [34]	FF++	-	64.10
(LQ) MATT [61]	FF++	99.80	67.02
(HQ) DSP-FWA [15]	FF++	-	69.30
(HQ) MTD-Net [36]	FF++	99.38	70.12
(HQ) Xception [13]	FF++	95.73	73.04
(HQ) CFFs [37]	FF++	97.03	74.20
(HQ) DAM [67]	FF++	-	75.30
(HQ) SOLA-sup [63]	FF++(DF)	-	76.02
(HQ) SPSL [20]	FF++	95.32	76.88
(HQ) LiSiam [38]	FF++	99.13	78.21
(HQ) SLADD [35]	FF++	98.40	79.70
(HQ) LipForensics [68]	FF++	99.70	82.40
(HQ) HFI-Net [39]	FF++	98.86	83.28
(HQ) TALL-ViT [73]	FF++	-	86.58
(HQ) TALL-Swin [73]	FF++	-	<b>90.79</b>
(HQ) FTCN [64]	FF++	-	86.90
(HQ) AltFreezing [72]	FF++	99.70	89.50
(HQ) F <sup>2</sup> Trans [40]	FF++	99.74	89.87
(RAW) PCL+I2G [17]	FF++	99.79	90.03
(HQ) Baseline	FF++	99.13	81.90
(HQ) <b>Baseline+BA (Ours)</b>	FF++	98.64	90.22
(LQ) F <sup>3</sup> -Net [19]	FF++	93.30	61.51
(LQ) RECCE [62]	FF++	95.02	68.71
(LQ) Two-branch [66]	FF++	93.18	73.41
(LQ) Baseline	FF++	94.25	75.68
(LQ) <b>Baseline+BA (Ours)</b>	FF++	91.24	<b>81.39</b>

TABLE V

PERFORMANCE COMPARISON OF OUR METHOD WITH STATE-OF-THE-ARTS UNDER CROSS-DATASETS SCENARIO FROM FF++ TO DFDCP AND FFIW

Method	Training Set	DFDCP [72]	FFIW [73]
PCL+I2G [17]	FF++(RAW)	74.37	-
SRM [21]	FF++(HQ)	79.70	-
SBI [18]	FF++(RAW)	86.15	<b>84.83</b>
SBI [18]	FF++(HQ)	85.51	83.22
Baseline	FF++(HQ)	80.37	63.85
<b>Baseline+BA (Ours)</b>	FF++(HQ)	<b>89.57</b>	76.76

TABLE VI

PERFORMANCE COMPARISON OF OUR METHOD WITH STATE-OF-THE-ARTS UNDER CROSS-FACESWAP&GAN SCENARIO FROM FF++ TO STYLEGAN

Method	DF	F2F	NT	FS	Avg.
(HQ) MATT [61]	43.99	43.99	43.99	43.99	43.99
(HQ) RECCE [62]	44.80	44.80	44.80	44.80	44.80
(HQ) SBI [18]	48.60	48.60	48.60	48.60	48.60
(HQ) <b>Baseline+BA (Ours)</b>	<b>68.55</b>	<b>58.06</b>	<b>55.10</b>	<b>59.08</b>	<b>60.20</b>
(LQ) MATT [61]	32.51	32.51	32.51	32.51	32.51
(LQ) RECCE [62]	58.95	58.95	58.95	58.95	58.95
(LQ) SBI [18]	64.11	<b>64.11</b>	64.11	64.11	64.11
(LQ) <b>Baseline+BA (Ours)</b>	<b>71.07</b>	52.51	<b>88.62</b>	<b>70.28</b>	<b>70.62</b>

DeepFakes, e.g., finding the blending artifacts, and they cannot handle StyleGAN synthesized faces. In contrast, our method

TABLE VII

EFFECT OF VARIOUS ADAPTATION SETTINGS

Method	DF → F2F		DF → FS		NT → FS	
	DF	F2F	DF	FS	NT	FS
Baseline	99.32	57.98	99.53	76.18	87.61	53.45
+ GRL [69]	99.29	68.32	99.55	89.46	88.22	71.11
+ MMD [44]	99.22	84.12	99.52	94.76	87.63	83.88
+ FA	99.30	86.49	99.62	95.61	87.90	84.53
+ SD	97.85	90.59	96.02	98.57	77.53	95.89
+ Ent. [70]	74.89	54.98	68.73	54.92	55.67	50.67
Baseline+BA (Ours)	96.43	92.13	98.08	98.09	85.03	93.94

outperforms these methods by a large margin, on average 37%, 34% and 24% better at HQ level, and 117%, 20% and 10% better at LQ level, than MATT, RECCE and SBI, respectively. This is because our method focuses on learning the common forgery knowledge in the central regions of both faceswap and GAN DeepFakes, instead of simply exposing the faceswap-specific forgery features.

### E. Ablation Study

We also conduct comprehensive ablation experiments to fully analyze the proposed method. More specifically, the ablation study 1) validates the efficacy of proposed approach; 2) demonstrates the proposed method is data-efficient and can work with various backbone networks; and 3) shows that our method can also benefit other augmentation-based face forgery detection methods to improve their generalization ability.

1) *Various Adaptation Settings*: To provide more insights, we further investigate the proposed bi-directional adaptation strategy on the LQ set of FF++. We use the DeepFake detector trained on the source domain without adaptation as the baseline. GRL [69] and MMD [44] are two classical domain adaptation methods. GRL uses the gradient reversal layer and MMD attempts to reduce the distance of the probability distributions between source and target domains. GRL is exactly the domain adaptation method adopted in FEAT [54]. We denote the forward adaptation method in our method as FA. In the top part of Table VII, we evaluate these three domain adaptation methods. It can be seen that adding GRL, MMD or FA can improve the performance in the target domain, and using our FA attains the highest performance gain.

Then we take a further step to analyze the effect of backward adaptation in the bottom part of Table VII, where SD means that the self-distillation is used in backward adaptation, while Ent. represents that self-distillation is replaced with entropy minimization of samples, which is inspired by the method of [70]. We observe that only adding SD can notably improve the performance on the target domain, but its performance on the source domain is compromised, especially in NT → FS, with 77.53% of +SD v.s. 87.61% of Baseline. This is because self-distillation is effective on the target domain, but overlooks the adaptation learned in forward adaptation. It can also be seen that replacing SD with Ent. is a bad idea, as this strategy degrades the detection performance in both domains. By adding both our forward adaptation and backward adaptation components to the baseline, our method can reach a






Fig. 7. Performance of using various amounts of training samples.

TABLE VIII  
EFFECT OF FREQUENCY MODULE

Method	DF → F2F		DF → FS		NT → FS	
	DF	F2F	DF	FS	NT	FS
C-192,D-4	93.34	90.36	96.56	97.51	80.97	92.69
C-384,D-4	95.26	90.65	97.82	97.66	84.13	93.46
C-768,D-2	95.17	90.97	96.97	97.35	84.21	92.39
C-768,D-3	95.70	91.21	97.86	96.87	84.23	92.78
C-768,D-4	96.43	92.13	98.08	98.09	85.03	93.94

good balance of the detection performance in both source and target domains. This clearly demonstrates that the proposed bi-directional adaptation is highly effective.

2) *Various Amounts of Training Samples*: This part investigates the data efficiency of the proposed method. We randomly select a proportion or percentage of the samples in target domain for training to simulate data-constrained scenarios. Fig. 7 depicts the performance of our method for the adaptation pairs of DF → F2F, DF → NT, and F2F → FS on the LQ set of FF++ using the percentage of the target samples ranging from 20% to 100%. As expected, increasing the percentage generally improves achievable performance. More importantly, noting the baseline+FA performance of 86.49 for DF → F2F from Table VII, it can be seen that our method can still considerably enhance the performance to over 90 even if only 20% of target samples are used. This indicates that our method can be utilized in the more restricted cases where the available target samples are insufficient.

3) *Effect of Frequency Module*: The number of channels (C) and depth (D) of DCT blocks are two important factors for extracting frequency features. The C and D values used in our method are 768 and 4. We now investigate the achievable performance of our method using various C and D values on the LQ set of FF++. As can be seen from Table VIII, given depth D = 4, increasing channels from C = 192, 384 to 768 improves the achievable performance on both source and target domains. Likewise, with channels fixed to C = 768, increasing depth from D = 2, 3 to 4 improves the achievable performance on both domains. The results of Table VIII also validate that our choice of C = 768 and D = 4 is appropriate.

4) *Various Feature Extractors*: As mentioned in the introduction section, the proposed framework is independent of the architecture. Table IX shows the performance of applying our framework on ResNet-50 [30], Xception [31], EfficientNet-B0 [32], EfficientNet-B4 [32], ViT-small [29] and ViT-base [29] on the LQ set of FF++. It can be seen that for each base network, our method significantly improves the performance on

TABLE IX  
PERFORMANCE OF USING VARIOUS FEATURE EXTRACTORS

Method	DF → F2F		DF → FS		NT → FS	
	DF	F2F	DF	FS	NT	FS
ResNet-50 [30]	99.40	57.51	99.40	65.72	87.72	57.27
ResNet-50 + Ours	96.89	76.17	96.49	90.61	77.99	71.21
Xception [31]	98.71	63.04	98.71	65.11	89.81	51.93
Xception + Ours	98.66	78.52	98.60	90.37	88.89	85.18
EfficientNet-B0 [32]	98.98	58.56	98.98	68.93	87.48	48.28
EfficientNet-B0 + Ours	90.26	91.24	98.32	94.66	84.45	91.57
EfficientNet-B4 [32]	98.82	57.91	98.82	68.63	88.77	48.56
EfficientNet-B4 + Ours	93.98	88.44	99.26	90.17	89.62	90.35
ViT-Small [29]	98.69	59.91	98.69	75.75	86.82	52.99
ViT-Small + Ours	95.63	87.46	97.27	95.58	83.83	83.17
ViT-Base [29]	99.32	57.98	99.53	76.18	87.61	53.45
ViT-Base + Ours	96.43	92.13	98.08	98.09	84.24	93.94

TABLE X  
PERFORMANCE OF ADDING OUR METHOD ON SBI

Adaptation Setting	SBI [18]	SBI + Ours
(HQ) SBI → DF	85.88	99.00
(HQ) SBI → F2F	80.07	97.23
(HQ) SBI → NT	72.47	90.28
(HQ) SBI → FS	75.62	92.92
Avg.	78.51	<b>94.86</b>
(LQ) SBI → DF	80.70	88.93
(LQ) SBI → F2F	67.63	68.20
(LQ) SBI → NT	65.13	64.98
(LQ) SBI → FS	62.50	65.52
Avg.	68.99	<b>71.91</b>

the target domain while maintaining a favorable performance on the source domain. This demonstrates that our proposed method is generically applicable.

5) *Added on State-of-the-Art DeepFake Detection*: This part investigates the effectiveness of our method to improve recently developed advanced detection methods. We use SBI [18] as an example. Specifically, we retrain SBI with HQ and LQ faces in FF++, and compare it with our method added on SBI. In this experiment, the source domain consists of fake faces blended by SBI and pristine faces. The target domain is each manipulation method. Thus, four new adaptation sub-tasks named SBI → DF, SBI → F2F, SBI → NT, and SBI → FS are formed. The results shown in Table X reveal that with our method added on, SBI notably improves the performance both at HQ and LQ levels. This indicates that our method can be easily integrated into other detection methods to enhance their generalization ability.

6) *Temperature  $\tau$* : We conduct experiments to investigate the influence of temperature in Eq. 4. As seen in Table XI, we observe that the model training with a large  $\tau$  (e.g. > 0.5) achieves stable performance. However, when  $\tau$  is set less than 0.5, the result will get unstable.

7) *Teacher Model Updating Strategy*: To validate the effectiveness of updating strategy, we conduct addition experiment using exponential moving average (EMA) updating

TABLE XI  
EFFECT OF TEMPERATURE  $\tau$

Sub-task	0.35	0.5	0.65	0.8
(HQ) DF $\rightarrow$ F2F	98.56	99.08	99.03	99.07
(HQ) FS $\rightarrow$ NT	94.43	96.27	95.60	94.25

TABLE XII  
EFFECT OF TEACHER MODEL UPDATING STRATEGY

Strategy	(HQ) DF $\rightarrow$ F2F	(HQ) FS $\rightarrow$ NT
EMA [77]	98.49	94.19
Baseline+BA (Ours)	99.08	96.27

TABLE XIII  
CROSS-DATASET GENERALIZATION COMPARISON

Adaptation Setting	CelebDF
FF++(HQ) $\rightarrow$ CelebDF	90.22
FF++(HQ) $\rightarrow$ DFDCP	85.24

strategy [77] on two sub-tasks (e.g. FS  $\rightarrow$  NT and DF  $\rightarrow$  F2F). The results are presented in Table XII. Our method achieves better performances than exponential moving average on both two sub-tasks. We conjecture this is because the direct assignment makes the knowledge transferred from teacher model to student model faster, which thereby conveys more comprehensive knowledge than EMA strategy.

8) *Results on Unseen Dataset:* We further conduct an experiment to evaluate the generalization performance on unseen dataset, where we make a comparison on two adaptation settings named FF++(HQ)  $\rightarrow$  CelebDF and FF++(HQ)  $\rightarrow$  DFDCP. As shown in Table XIII, We observe that our method can also achieve comparable generalization performance on Celeb-DF when the target domain is DFDCP.

## V. CONCLUSION

This paper has proposed a new framework to detect new forgeries across different domains, called *DomainForensics*. Different from the recent methods, which empirically seeks the common traces on data view, our method aims to transfer the forgery knowledge from known forgeries (fully labeled source domain) to new forgeries (label-free target domain). Since the general domain adaptation methods are not competent to capture the forgery features, we have designed a new bi-directional adaptation strategy that considers both the forward adaptation and backward adaptation. Specifically, the forward adaptation transfers the knowledge from the source to the target domain, and the backward adaptation reverses the adaptation from the target to the source domain. With this backward adaptation, the detector can be further enhanced to learn new forgery features from unlabeled data and avoid forgetting the known knowledge of known forgery. Extensive experiments have been conducted on three datasets with the comparison to several existing state-of-the-art counterparts. The results obtained have demonstrated that our method is effective in exposing new forgeries, and it can be integrated into various architectures to improve their generalization ability.

## REFERENCES

- [1] I. J. Goodfellow et al., “Generative adversarial nets,” in *Proc. NIPS*, 2014, pp. 1–9.
- [2] J. Thies, M. Zollhöfer, M. Stamminger, C. Theobalt, and M. Nießner, “Face2Face: Real-time face capture and reenactment of RGB videos,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2387–2395.
- [3] S. Suwajanakorn, S. M. Seitz, and I. Kemelmacher-Shlizerman, “Synthesizing Obama: Learning lip sync from audio,” *ACM Trans. Graph.*, vol. 36, no. 4, pp. 1–13, Aug. 2017.
- [4] H. Kim et al., “Deep video portraits,” *ACM Trans. Graph.*, vol. 37, no. 4, pp. 1–14, Aug. 2018.
- [5] T. Karras, T. Aila, S. Laine, and J. Lehtinen, “Progressive growing of GANs for improved quality, stability, and variation,” in *Proc. Int. Conf. Learn. Represent.*, 2018, pp. 1–26.
- [6] T. Karras, S. Laine, and T. Aila, “A style-based generator architecture for generative adversarial networks,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 4401–4410.
- [7] B. Chesney and D. Citron, “Deep fakes: A looming challenge for privacy, democracy, and national security,” *Calif. Law Rev.*, vol. 107, no. 6, pp. 1753–1820, 2019.
- [8] Y. Li, M.-C. Chang, and S. Lyu, “In ictu oculi: Exposing AI created fake videos by detecting eye blinking,” in *Proc. IEEE Int. Workshop Inf. Forensics Secur. (WIFS)*, Hong Kong, Dec. 2018, pp. 1–7.
- [9] D. Afchar, V. Nozick, J. Yamagishi, and I. Echizen, “MesoNet: A compact facial video forgery detection network,” in *Proc. IEEE Int. Workshop Inf. Forensics Secur. (WIFS)*, Dec. 2018, pp. 1–7.
- [10] J. Li, H. Xie, J. Li, Z. Wang, and Y. Zhang, “Frequency-aware discriminative feature learning supervised by single-center loss for face forgery detection,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 6458–6467.
- [11] Z. Sun, Y. Han, Z. Hua, N. Ruan, and W. Jia, “Improving the efficiency and robustness of DeepFakes detection through precise geometric features,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 3609–3618.
- [12] Y. He et al., “ForgeryNet: A versatile benchmark for comprehensive forgery analysis,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 4358–4367.
- [13] A. Rossler, D. Cozzolino, L. Verdoliva, C. Riess, J. Thies, and M. Niessner, “FaceForensics++: Learning to detect manipulated facial images,” in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 1–11.
- [14] Y. Li, X. Yang, P. Sun, H. Qi, and S. Lyu, “Celeb-DF: A large-scale challenging dataset for DeepFake forensics,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 3207–3216.
- [15] Y. Li and S. Lyu, “Exposing deepfake videos by detecting face warping artifacts,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, Jun. 2019, pp. 46–52.
- [16] L. Li et al., “Face X-ray for more general face forgery detection,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 5001–5010.
- [17] T. Zhao, X. Xu, M. Xu, H. Ding, Y. Xiong, and W. Xia, “Learning self-consistency for deepfake detection,” in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, Sep. 2021, pp. 15023–15033.
- [18] K. Shiohara and T. Yamasaki, “Detecting deepfakes with self-blended images,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, New Orleans, LA, USA, Jun. 2022, pp. 18699–18708.
- [19] Y. Qian et al., “Thinking in frequency: Face forgery detection by mining frequency-aware clues,” in *Proc. ECCV*, Aug. 2020, pp. 1–18.
- [20] H. Liu et al., “Spatial-phase shallow learning: Rethinking face forgery detection in frequency domain,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 772–781.
- [21] Y. Luo, Y. Zhang, J. Yan, and W. Liu, “Generalizing face forgery detection with high-frequency features,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 16317–16326.
- [22] C. Dong, A. Kumar, and E. Liu, “Think twice before detecting GAN-generated fake images from their spectral domain imprints,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 7855–7864.
- [23] D. Cozzolino, J. Thies, A. Rössler, C. Riess, M. Nießner, and L. Verdoliva, “ForensicTransfer: Weakly-supervised domain adaptation for forgery detection,” 2018, *arXiv:1812.02510*.
- [24] S. Aneja and M. Nießner, “Generalized zero and few-shot transfer for facial forgery detection,” 2020, *arXiv:2006.11863*.

- [25] H. Qiu et al., "Few-shot forgery detection via guided adversarial interpolation," *Pattern Recognit.*, vol. 144, Dec. 2023, Art. no. 109863.
- [26] H. Wang, X. Guo, Z. Deng, and Y. Lu, "Rethinking minimal sufficient representation in contrastive learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 16041–16050.
- [27] T. Chen et al., "Big self-supervised models are strong semi-supervised learners," in *Proc. NeurIPS*, 2020, pp. 1–13.
- [28] Z. Li, R. Cai, H. Li, K. Lam, Y. Hu, and A. C. Kot, "One-class knowledge distillation for face presentation attack detection," *IEEE Trans. Inf. Forensics Security*, vol. 17, pp. 2137–2150, 2022.
- [29] A. Dosovitskiy et al., "An image is worth  $16 \times 16$  words: Transformers for image recognition at scale," in *Proc. ICLR*, 2021, pp. 1–21.
- [30] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [31] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1251–1258.
- [32] M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 1–11.
- [33] S. Agarwal, H. Farid, Y. Gu, M. He, K. Nagano, and H. Li, "Protecting world leaders against deep fakes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, Jun. 2019, pp. 38–45.
- [34] K. Sun et al., "Domain general face forgery detection by learning to weight," in *Proc. AAAI Conf. Artif. Intell.*, vol. 35, 2021, pp. 2638–2646.
- [35] L. Chen, Y. Zhang, Y. Song, L. Liu, and J. Wang, "Self-supervised learning of adversarial example: Towards good generalizations for deepfake detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, New Orleans, LA, USA, Jun. 2022, pp. 18689–18698.
- [36] J. Yang, A. Li, S. Xiao, W. Lu, and X. Gao, "MTD-Net: Learning to detect deepfakes images by multi-scale texture difference," *IEEE Trans. Inf. Forensics Security*, vol. 16, pp. 4234–4245, 2021.
- [37] P. Yu, J. Fei, Z. Xia, Z. Zhou, and J. Weng, "Improving generalization by commonality learning in face forgery detection," *IEEE Trans. Inf. Forensics Security*, vol. 17, pp. 547–558, 2022.
- [38] J. Wang, Y. Sun, and J. Tang, "LiSiam: Localization invariance Siamese network for deepfake detection," *IEEE Trans. Inf. Forensics Security*, vol. 17, pp. 2425–2436, 2022.
- [39] C. Miao, Z. Tan, Q. Chu, N. Yu, and G. Guo, "Hierarchical frequency-assisted interactive networks for face manipulation detection," *IEEE Trans. Inf. Forensics Security*, vol. 17, pp. 3008–3021, 2022.
- [40] C. Miao, Z. Tan, Q. Chu, H. Liu, H. Hu, and N. Yu, "F2Trans: High-frequency fine-grained transformer for face forgery detection," *IEEE Trans. Inf. Forensics Security*, vol. 18, pp. 1039–1051, 2023.
- [41] X. Zhu, H. Wang, H. Fei, Z. Lei, and S. Z. Li, "Face forgery detection by 3D decomposition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 2928–2938.
- [42] S. Ben-David, J. Blitzer, K. Crammer, and F. Pereira, "Analysis of representations for domain adaptation," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, vol. 19, 2006, pp. 1–8.
- [43] E. Tzeng, J. Hoffman, N. Zhang, K. Saenko, and T. Darrell, "Deep domain confusion: Maximizing for domain invariance," 2014, *arXiv:1412.3474*.
- [44] M. Long, Y. Cao, J. Wang, and M. Jordan, "Learning transferable features with deep adaptation networks," in *Proc. 32nd Int. Conf. Mach. Learn.*, vol. 37, Jul. 2015, pp. 97–105.
- [45] B. Sun and K. Saenko, "Deep CORAL: Correlation alignment for deep domain adaptation," in *Proc. ECCV*, Oct. 2016, pp. 443–450.
- [46] Y. Zhang, T. Liu, M. Long, and M. Jordan, "Bridging theory and algorithm for domain adaptation," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, pp. 7404–7413.
- [47] I. J. Goodfellow et al., "Generative adversarial networks," *Commun. ACM*, vol. 63, no. 11, pp. 139–144, 2020.
- [48] Y. Ganin et al., "Domain-adversarial training of neural networks," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 2030–2096, May 2016.
- [49] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial discriminative domain adaptation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 7167–7176.
- [50] M. Xu et al., "Adversarial domain adaptation with domain mixup," in *Proc. AAAI Conf. Artif. Intell.*, 2020, vol. 34, no. 4, pp. 6502–6509.
- [51] J. Yang, J. Liu, N. Xu, and J. Huang, "TVT: Transferable vision transformer for unsupervised domain adaptation," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis. (WACV)*, Jan. 2023, pp. 520–530.
- [52] T. Xu, W. Chen, P. Wang, F. Wang, H. Li, and R. Jin, "CDTrans: Cross-domain transformer for unsupervised domain adaptation," 2021, *arXiv:2109.06165*.
- [53] T. Sun, C. Lu, T. Zhang, and H. Ling, "Safe self-refinement for transformer-based domain adaptation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 7191–7200.
- [54] B. Chen and S. Tan, "FeatureTransfer: Unsupervised domain adaptation for cross-domain deepfake detection," *Secur. Commun. Netw.*, vol. 2021, pp. 1–8, Jun. 2021.
- [55] Y. Li, X. Bian, M.-C. Chang, and S. Lyu, "Exploring the vulnerability of single shot module in object detectors via imperceptible background patches," in *Proc. BMVC*, Sep. 2019, pp. 1–12.
- [56] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 618–626.
- [57] A. Vaswani, "Attention is all you need," in *Proc. 31st Conf. Neural Inf. Process. Syst.*, Long Beach, CA, USA, 2017, pp. 1–11.
- [58] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 12, pp. 4217–4228, Dec. 2021.
- [59] J. Deng, J. Guo, E. Ververas, I. Kotsia, and S. Zafeiriou, "RetinaFace: Single-shot multi-level face localisation in the wild," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 5202–5211.
- [60] A. Paszke et al., "Automatic differentiation in pytorch," in *Proc. NIPS*, Long Beach, CA, USA, Dec. 2017, pp. 1–4.
- [61] H. Zhao, T. Wei, W. Zhou, W. Zhang, D. Chen, and N. Yu, "Multi-attentional deepfake detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 2185–2194.
- [62] J. Cao, C. Ma, T. Yao, S. Chen, S. Ding, and X. Yang, "End-to-end reconstruction-classification learning for face forgery detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 4113–4122.
- [63] J. Fei, Y. Dai, P. Yu, T. Shen, Z. Xia, and J. Weng, "Learning second order local anomaly for general face forgery detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, New Orleans, LA, USA, Jun. 2022, pp. 20238–20248.
- [64] Y. Zheng, J. Bao, D. Chen, M. Zeng, and F. Wen, "Exploring temporal coherence for more general video face forgery detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 15044–15054.
- [65] L. Van der Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, no. 11, pp. 2579–2605, 2008.
- [66] I. Masi, A. Killekar, and R. M. Mascarenhas, "Two-branch recurrent network for isolating deepfakes in videos," in *Proc. ECCV*, Glasgow, U.K., Aug. 2020, pp. 667–684.
- [67] T. Zhou, W. Wang, Z. Liang, and J. Shen, "Face forensics in the wild," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2021, pp. 5778–5788.
- [68] A. Haliassos, K. Vougioukas, S. Petridis, and M. Pantic, "Lips don't lie: A generalisable and robust approach to face forgery detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 5039–5049.
- [69] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," in *Proc. ICML*, Lille, France, 2015, pp. 1180–1189.
- [70] Y. Grandvalet and Y. Bengio, "Semi-supervised learning by entropy minimization," in *Proc. 17th Int. Conf. Neural Inf. Process. Syst.*, 2004, pp. 529–536.
- [71] B. Dolhansky, R. Howes, B. Pfau, N. Baram, and C. C. Ferrer, "The deepfake detection challenge (DFDC) preview dataset," 2019, *arXiv:1910.08854*.
- [72] Z. Wang, J. Bao, W. Zhou, W. Wang, and H. Li, "AltFreezing for more general video face forgery detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2023, pp. 4129–4138.
- [73] Y. Xu, J. Liang, G. Jia, Z. Yang, Y. Zhang, and R. He, "TALL: Thumbnail layout for deepfake video detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2023, pp. 22601–22611.
- [74] Z. Wang et al., "DIRE for diffusion-generated image detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2023, pp. 22388–22398.
- [75] D. Lee, "Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks," *Proc. ICML* 2013, pp. 1–6.



- [76] I. Zeki Yalniz, H. Jégou, K. Chen, M. Paluri, and D. Mahajan, “Billion-scale semi-supervised learning for image classification,” 2019, *arXiv:1905.00546*.
- [77] A. Tarvainen and H. Valpola, “Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results,” in *Proc. Int. Conf. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 1195–1204.



**Qingxuan Lv** was born in Shanxi, China, in 1996. He received the bachelor's degree in computer science and technology from Shanxi University of Finance and Economics in 2018 and the master's degree in computer science and technology from the Ocean University of China (OUC) in 2021, where he is currently pursuing the Ph.D. degree with the Ocean Group, VisionLab. His research interests include computer vision and machine learning. Specifically, he is interested in universal domain adaptation and semantic segmentation.



**Yuezun Li** (Member, IEEE) received the B.S. degree in software engineering from Shandong University in 2012, the M.S. degree in computer science in 2015, and the Ph.D. degree in computer science from University at Albany, SUNY, in 2020. He was a Senior Research Scientist with the Department of Computer Science and Engineering, University at Buffalo, SUNY. He is currently a Lecturer with the Center on Artificial Intelligence, Ocean University of China. His work has been published in peer-reviewed conferences and journals, including ICCV, CVPR, IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY, IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY. His research interests include computer vision and multimedia forensics.



**Junyu Dong** received the B.Sc. and M.Sc. degrees in applied mathematics from the Department of Applied Mathematics, Ocean University of China, Qingdao, China, in 1993 and 1999, respectively, and the Ph.D. degree in image processing from the Department of Computer Science, Heriot-Watt University, Edinburgh, U.K., in November 2003. He is currently a Professor and the Head of the Department of Computer Science and Technology. His research interests include machine learning, big data, computer vision, and underwater image processing.



**Sheng Chen** (Life Fellow, IEEE) received the B.Eng. degree in control engineering from East China Petroleum Institute, Dongying, China, in 1982, the Ph.D. degree in control engineering from the City, University of London, in 1986, and the D.Sc. degree from the University of Southampton, Southampton, U.K., in 2005. From 1986 to 1999, he held research and academic appointments with The University of Sheffield, The University of Edinburgh, and the University of Portsmouth, U.K. Since 1999, he has been with the School of Electronics and Computer Science, University of Southampton, where he is currently a Professor of intelligent systems and signal processing. He has authored over 700 research articles. He has more than 20,000 Web of Science citations with an H-index of 61 and more than 39,000 Google Scholar citations with an H-index of 83. His research interests include adaptive signal processing, wireless communications, modeling and identification of nonlinear systems, neural network and machine learning, intelligent control system design, evolutionary computation methods, and optimization. He was elected to a fellow of the United Kingdom Royal Academy of Engineering in 2014. He is a fellow of Asia-Pacific Artificial Intelligence Association (FAAIA) and IET and an original ISI highly cited researcher in engineering in March 2004.



**Hui Yu** received the Ph.D. degree from Brunel University London in 2009. He is currently a Professor of visual computing with the University of Portsmouth, U.K., and leads the Visual Computing Group. Before joining the University of Portsmouth in 2012, he was with the University of Glasgow and Queen's University Belfast. His research interests include visual computing and machine learning with application to 3D/4D facial expression tracking, reconstruction and affective analysis, human-machine and social interaction, virtual and augmented reality, and intelligent vehicles. He serves as an Associate Editor for IEEE TRANSACTIONS ON HUMAN-MACHINE SYSTEMS, IEEE TRANSACTIONS ON INTELLIGENT VEHICLES, and IEEE TRANSACTIONS ON COMPUTATIONAL SOCIAL SYSTEMS journal.



**Huiyu Zhou** received the B.Eng. degree in radio technology from Huazhong University of Science and Technology, Wuhan, China, in 1990, the M.Sc. degree in biomedical engineering from the University of Dundee, Dundee, U.K., in 2002, and the Ph.D. degree in computer vision from Heriot-Watt University, Edinburgh, U.K., in 2006. He is currently a Full Professor with the School of Computing and Mathematical Sciences, University of Leicester, Leicester, U.K. He has authored or co-authored over 450 peer-reviewed articles in the field.



**Shu Zhang** received the Ph.D. degree in computer application technologies from the Ocean University of China, Qingdao, China. He is currently an Associate Professor and a Postgraduate Supervisor with the Ocean University of China. Previously, he was a Research Associate with the University of Portsmouth, Portsmouth, U.K. His research interests include computer vision, feature analysis, 3D reconstruction, video processing, underwater image analysis, and deep learning among others.