

# Lips Don't Lie: A Generalisable and Robust Approach to Face Forgery Detection

Alexandros Haliassos<sup>1,†</sup> Konstantinos Vougioukas<sup>1</sup> Stavros Petridis<sup>1</sup> Maja Pantic<sup>1,2</sup>

<sup>1</sup>Imperial College London <sup>2</sup>Facebook London

{alexandros.haliassos14,k.vougioukas,stavros.petridis04,m.pantic}@imperial.ac.uk

## Abstract

Although current deep learning-based face forgery detectors achieve impressive performance in constrained scenarios, they are vulnerable to samples created by unseen manipulation methods. Some recent works show improvements in generalisation but rely on cues that are easily corrupted by common post-processing operations such as compression. In this paper, we propose *LipForensics*, a detection approach capable of both generalising to novel manipulations and withstanding various distortions. *LipForensics* targets high-level semantic irregularities in mouth movements, which are common in many generated videos. It consists in first pretraining a spatio-temporal network to perform visual speech recognition (lipreading), thus learning rich internal representations related to natural mouth motion. A temporal network is subsequently finetuned on fixed mouth embeddings of real and forged data in order to detect fake videos based on mouth movements without overfitting to low-level, manipulation-specific artefacts. Extensive experiments show that this simple approach significantly surpasses the state-of-the-art in terms of generalisation to unseen manipulations and robustness to perturbations, as well as shed light on the factors responsible for its performance.

## 1. Introduction

Recent advances in deep generative models, especially Generative Adversarial Networks (GANs) [24], have enabled the use of off-the-shelf models that can produce ultra-realistic fake videos with little human effort or expertise. Face manipulation methods in particular have raised considerable concerns due to their ability to alter a person's identity [38, 30, 34], expression [52, 51, 44], or lip movements [50, 57, 56] to match the face in a given target video. The misuse of such technology can spread political propaganda, defame individuals, or damage trust in journalism.

With the aid of recent releases of large-scale face forgery

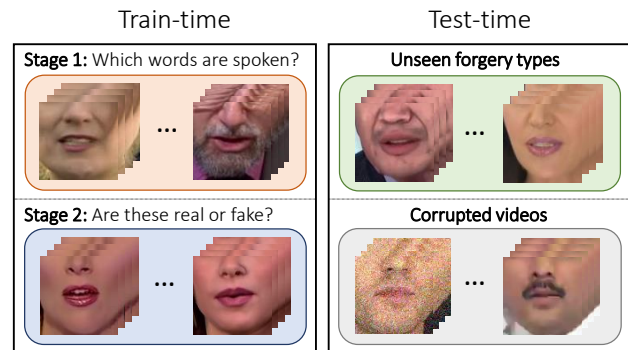


Figure 1. **Illustration of our training and testing objectives.** By leveraging semantically high-level, spatio-temporal representations learned via the task of lipreading, our face forgery detector can handle both novel manipulations and common corruptions at test-time.

datasets [46, 38, 30, 16], it is possible to train deep convolutional neural networks (CNNs) to detect fake videos [3, 25, 64, 4, 14, 45, 42, 36]. Despite excellent accuracy on samples that are independent and identically distributed to the training distribution, it is known that dramatic drops in performance may be experienced when this stringent criterion is not met [35, 8, 43]. For example, detectors often significantly underperform on novel forgery types. This understandably sparks concerns, as a deployed detector is unlikely to be exposed only to forgeries seen during training. Moreover, they are often sensitive to common perturbations such as compression and are, as a result, vulnerable to the image processing pipelines on social networks.

Recent attempts to boost generalisation to novel forgeries include simple data augmentations [59], a two-branch network that amplifies multi-band frequencies [41], an autoencoder-like structure to serve as an anomaly detector [13, 17, 43], and patch-based classification to model local patterns [8]. However, these methods still substantially overfit to seen manipulations. A particularly effective method is Face X-ray [35], which proposes to predict the blending boundary between the background image and the inserted, altered face. Although it attains impressive gen-

<sup>†</sup>Corresponding author.

eralisation in cross-manipulation experiments, it relies on often imperceptible patterns which are susceptible to low-level post-processing operations.

It is natural to ask: Are there *semantically high-level* inconsistencies across manipulation methods, which are thus more robust to routine perturbations? We observe that most face video forgeries alter the mouth in some way, to match it with someone else’s identity, speech, or expression. Due to the intricate motion of the mouth, current manipulation methods find it difficult to generate movements without falling into the “uncanny valley<sup>1</sup>.” For example, fake mouths often do not adequately close when they pronounce certain phonemes [4]. We also notice unnatural fluctuations in the speed of movements, as well as alterations in the shape of the mouth or its interior (*e.g.*, teeth) from frame-to-frame, even when there is no speech (see supplementary material for examples). These irregularities provide a precious opportunity for detectors to capitalise on; yet, to the best of our knowledge, previous works do not specifically target mouth motion using spatio-temporal neural networks.

The human visual system can perceive subtle anomalies in forged oral movements as a result of extensive experience in observing real mouths move. In order to endow a network with such experience, we propose to pretrain it on a large corpus of real videos to perform the difficult task of visual speech recognition, also known as lipreading. To be capable of disambiguating similar words, the network must learn rich spatio-temporal representations related to the mouth as well as the teeth and tongue [12]. Next, we transfer the acquired knowledge to face forgery detection. Crucially, we treat the first part of the network as a frozen feature extractor that outputs an embedding per frame, and only finetune the temporal convolutional network that takes these embeddings as input. This prevents the network from learning to discriminate the data based on unstable, low-level patterns that may not be generated by other forgery methods.

We dub our approach *LipForensics* (see Figure 1). We conduct extensive experiments to compare its performance with the state-of-the-art in various challenging scenarios. We find that, in most cases, it significantly outperforms previous methods with respect to generalisation to unseen forgeries, while exhibiting remarkable robustness to common corruptions which degrade other models’ performance. Further, in-distribution experiments reveal that *LipForensics* can effectively learn even on heavily compressed data, unlike other detectors. Finally, we validate our design choices through ablation studies, and compare with other large-scale pretraining tasks to demonstrate the superiority of lipreading for achieving generalisable and robust face forgery detection.

---

<sup>1</sup>The “uncanny valley” refers to the unease experienced by humans when observing a realistic computer-generated face.

## 2. Related Work

**Face forgery detection.** Some earlier face forgery detection works bias the network away from learning high-level features by constraining CNN filters [6] or using relatively shallow networks [3] to focus on mesoscopic features. Rössler *et al.* [46] showed, however, that these are outperformed by a deep, unconstrained Xception [10] network. Some methods [5, 49] predict the manipulated region along with the label in a multi-task fashion. Another line of research relates to the observation that fully CNN-generated images may exhibit anomalous patterns in the frequency spectrum [19, 18, 63, 45]. In [25, 47], a CNN followed by an LSTM [27] are used to capture spatio-temporal features. In [58], fake frames are detected based on anomalies in the neuron behaviour of a face recognition network.

Some works exploit the correspondence between the visual and auditory modalities [4, 42, 11, 33]. For example, [4] targets inconsistencies between the “M”, “B”, and “P” phonemes and their corresponding visemes, requiring voice-to-text transcriptions and viseme-phoneme alignment. We note that, in contrast, our method is visual-only, *i.e.*, does not take as input the audio during training or inference. A concurrent work [61] also exploits lip movements but is a biometric approach, hence requiring a reference video at inference.

**Generalisation to novel manipulations.** Although current detectors tend to perform well when the training and test distributions are similar, cross-dataset experiments reveal that they underperform on unseen forgeries [35, 8, 38, 43]. To address this issue, ForensicTransfer [13] proposes an autoencoder-like structure to prevent the network from discarding frame information. Follow-up works LAE [17] and MTDS [43] additionally exploit binary forgery segmentation masks. In [41], a two-branch recurrent neural network amplifies multi-band frequencies. Wang *et al.* [59] apply blurring and compression augmentations, which help generalisation across fully-synthetic fake images. Chai *et al.* [8] instead hypothesise that local patterns generalise better and truncate image classifiers to reduce their receptive fields. We challenge this hypothesis by showing that high-level temporal inconsistencies in the mouth region can generalise very well (and lead to more robust detectors).

A particularly successful self-supervised approach is to generate fake videos on-the-fly that contain only the specific irregularities that one may want to target. For instance, FWA [37] targets the common affine face warping step in forgery pipelines. More recently, Li *et al.* [35] observe that many forgery algorithms depend on a blending step to realistically insert an altered face into a background image. They predict the blending boundary between the two images to achieve the current state-of-the-art generalisation

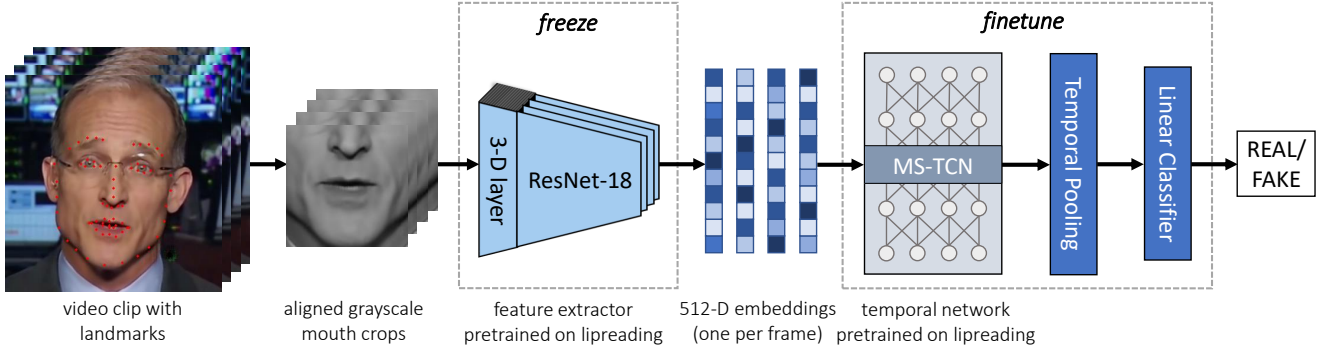


Figure 2. **Overview of the finetuning phase on face forgery detection.** The input to the network consists of 25 grayscale, aligned mouth crops (we only show four for illustration purposes). They pass through a frozen feature extractor (a ResNet-18 with an initial 3-D convolutional layer), which has been pretrained on lipreading and hence outputs embeddings sensitive to mouth movements. A multi-scale temporal convolutional network (MS-TCN), also pretrained on lipreading, is finetuned to detect fake videos based on semantically high-level irregularities in mouth motion.

on cross-dataset experiments. However, the artefacts that the network focuses on are susceptible to common perturbations, rendering it unsuitable for many real-life scenarios.

### 3. LipForensics

#### 3.1. Overview

We tackle the problem of building a generalisable and robust face forgery classifier by distinguishing between natural and anomalous mouth movements. We hypothesise that irregularities in mouth motion exist in fake videos regardless of the generation method that produced them. Due to the semantically high-level nature of these cues, they are also less easily corrupted by common perturbations. However, a spatio-temporal CNN simply trained on mouth crops will not necessarily learn the desired features, as it may separate the data based on other more manipulation-specific cues. To account for this, we take a two-step approach.

First, we pretrain a CNN, consisting of a spatio-temporal feature extractor followed by a temporal convolutional network, on the task of lipreading. We expect this process to induce internal representations that are sensitive to anomalous dynamics of the mouth in a high-level semantic space, since low-level patterns are likely insufficient for solving the task. Such pretraining is consistent with recent anomaly detection literature, which suggests that training on the “normal” class (in this case real videos) for a suitable task promotes learning of features that are useful for detecting “anomalous” samples (in this case fake videos) [23, 7].

Second, we freeze the feature extractor and finetune only the temporal network on forgery data; otherwise, the network may still learn to rely on unwanted artefacts rather than mouth movements. Other works attempt to alleviate overfitting to unstable cues by blurring or adding noise to the input [60, 59], forcing the network to target blending artefacts [35], or discentivising the model from discarding

frame information via a reconstruction loss [13, 17, 43]. In contrast, we pass video clips through a deep feature extractor that was trained to perform lipreading and, as a result, outputs embeddings that are relatively invariant to low-level artefacts. We illustrate this process in Figure 2.

#### 3.2. Formulation

Let  $\mathcal{X}$  be the set of grayscale video clips, real or fake, centred around the mouth, and let  $\mathcal{X}_r \subset \mathcal{X}$  denote the set of real clips. We are given a face forgery dataset  $\{(x_f^j, y_f^j)\}_{j=1}^{N_f}$  of size  $N_f$ , where  $x_f^j \in \mathcal{X}$  is a video clip and  $y_f^j \in \{0, 1\}$  denotes whether the clip is real or fake. We further assume that we are provided with a labelled lipreading dataset  $\{(x_l^j, y_l^j)\}_{j=1}^{N_l}$  of size  $N_l$ , where  $x_l^j \in \mathcal{X}_r$  is a real video clip of a word utterance and  $y_l^j \in \{0, \dots, L-1\}$  is a label specifying which word was spoken from a vocabulary of length  $L$ .

We wish to first train a multi-class neural network,  $f_l$ , to perform lipreading on real videos. It comprises a spatio-temporal feature extractor, a temporal net, and a linear classifier, parameterised by  $\theta_g$ ,  $\theta_h$ , and  $\theta_l$ , respectively. The parameters of these subnetworks are randomly initialised and then optimised together to minimise the standard cross entropy loss,  $\frac{1}{N_l} \sum_{j=1}^{N_l} \mathcal{L}_{\text{CE}}(f_l(x_l^j), y_l^j; \theta_g, \theta_h, \theta_l)$ .

To learn our forgery detector,  $f_f$ , we transfer  $\theta_g$  and  $\theta_h$  and replace the classifier with a binary one, parameterised by  $\theta_f$ . The temporal net is finetuned (and the classifier is trained from scratch) to minimise the binary cross entropy loss,  $\frac{1}{N_f} \sum_{j=1}^{N_f} \mathcal{L}_{\text{BCE}}(f_f(x_f^j), y_f^j; \theta_h, \theta_f)$ . The feature extractor is kept fixed during this phase.

**Architecture.** The feature extractor is a ResNet-18 [26] with an initial 3-D convolutional layer which preserves the temporal dimension via padding. The feature extractor outputs a 512-D vector for each input frame. The temporal

net is a multi-scale temporal convolutional network (MS-TCN) [40], combining short- and long-term temporal information at every layer by concatenating outputs of multiple branches, each with a different temporal receptive field. After a temporal global average pooling layer, a task-specific linear classifier outputs the estimated class probabilities. Architecture details are in the supplementary material.

**Lipreading pretraining.** The model is pretrained on Lipreading in the Wild (LRW) [12], a dataset containing over 500,000 utterances spanning hundreds of speakers in various poses. It is trained using the approach proposed in [39], which employs born-again distillation [20]. We use the student classifier in the third generation of teacher-student training. Unless stated otherwise, we use the publicly available, pretrained model found here<sup>2</sup>.

**Preprocessing.** The faces are detected and then aligned to the mean face, after which a scaled  $96 \times 96$  region around the mouth is cropped and transformed to grayscale. Our clips comprise 25 frames; thus, following random cropping, our input tensor is of size  $25 \times 88 \times 88 \times 1$ . Note that this size corresponds to a similar number of entries as the standard RGB frame input to many of the forgery detectors in the literature (commonly of size  $1 \times 256 \times 256 \times 3$ ) [35]. More preprocessing details are in the supplementary material, where we also consider the effect of using the full face as input rather than mouth crops.

**Training the forgery detector.** We use a batch size of 32 and Adam [32] optimisation with a learning rate of  $2 \times 10^{-4}$ . To address any data imbalance, we oversample the minority class. We terminate training when there is negligible improvement to the validation loss for 10 epochs. As data augmentation, we randomly crop the clips with size  $88 \times 88$  and horizontally flip with probability 0.5.

## 4. Experiments

### 4.1. Setup

**Datasets.** As in [35, 41], we mostly use **FaceForensics++** (FF++) [46] as our training dataset due to its forgery diversity. It contains 1.8 million manipulated frames and 4,000 fake videos generated using two face swapping algorithms, DeepFakes [1] and FaceSwap [2], and two face reenactment methods, Face2Face [52] and NeuralTextures [51]. As recommended in [46], we only use the first 270 frames for each training video, and the first 110 frames for each validation/testing video. Other datasets in our experiments include **DeeperForensics** [30] and **FaceShifter** [34], each featuring an improved face swapping algorithm applied to the real videos from FF++. Further, we use the test set of the

<sup>2</sup>[https://github.com/mpc001/Lipreading\\_using\\_Temporal\\_Convolutional\\_Networks](https://github.com/mpc001/Lipreading_using_Temporal_Convolutional_Networks)

face swapping **Celeb-DF-v2** [38] dataset. Finally, we use 3,215 test set videos from the **DeepFake Detection Challenge** (DFDC) [16], where subjects were filmed in extreme conditions, such as large poses and low lighting. More information is in the supplementary material.

**Metrics.** We report results using accuracy (following [46]) and/or area under the receiver operating characteristic curve (AUC; following [38, 35]). Since most existing models use a single frame as input, we compute video-level measures for fair comparison: we average the model predictions (each prediction corresponding to a frame or a video clip) across the entire video, as in [41]. As a result, *all models' video predictions are based on the same number of frames.*

### 4.2. Generalisation to unseen manipulations

Given the rapid advancements in forgery generation, we desire a detector that can correctly classify samples from novel manipulation methods after it is trained on a set of known forgeries. We simulate this scenario below.

**Models for comparison.** For comparison on our experiments using video-level metrics, we train various state-of-the-art detectors designed to improve cross-dataset generalisation as well as some other popular baselines, including: (1) **Face X-ray** [35]: we train an HRNet-W48 [48] both with constructed blended images and fake samples from the considered datasets. (2) **CNN-aug** [59]: we employ a ResNet-50 with JPEG compression and Gaussian blurring augmentations, both with probability of 0.1. (3) **Patch-based** [8]: we train an Xception classifier truncated after block 3 and average the predictions across the patches. We align the frames to remove rotation variation. (4) **Xception** [46]: we train the popular Xception baseline. (5) **CNN-GRU** [47]: we train a DenseNet-161 [29] followed by a GRU [9], to compare with a temporal model. All these models take as input a single RGB frame of the full face, except for CNN-GRU, which takes a 5-frame RGB clip. All but Patch-based use ImageNet pretrained weights. See supplementary material for more details. We also evaluate the effects of (1) training our spatio-temporal detector from random initialisation (“Scratch”), (2) finetuning the whole network after pretraining on LRW (“Ft whole”), and (3) finetuning only the temporal net (“LipForensics”), all of which are trained on grayscale mouth crops.

**Cross-manipulation generalisation.** To directly assess generalisation to unseen manipulations, without introducing confounders such as variations in pose or illumination, we experiment here with fake data that were created from the same source videos. Specifically, we test on each of the four methods in FF++ after we train on the remaining three. Rather than using raw videos (as in [35]), we use the high quality (HQ) subset of the dataset (as in *e.g.*, [41, 8, 43]),

Method	Train on remaining three				Avg
	DF	FS	F2F	NT	
Xception [46]	93.9	51.2	86.8	79.7	77.9
CNN-aug [59]	87.5	56.3	80.1	67.8	72.9
Patch-based [8]	94.0	60.5	87.3	84.8	81.7
Face X-ray [35]	99.5	<b>93.2</b>	94.5	92.5	94.9
CNN-GRU [47]	97.6	47.6	85.8	86.6	79.4
Scratch	93.0	56.7	98.8	98.3	86.7
Ft whole	98.4	80.4	99.4	<b>99.3</b>	94.4
LipForensics (ours)	<b>99.7</b>	90.1	<b>99.7</b>	99.1	<b>97.1</b>

Table 1. **Cross-manipulation generalisation.** Video-level AUC (%) when testing on each forgery type of FaceForensics++ HQ after training on the remaining three. The types are Deepfakes (DF), FaceSwap (FS), Face2Face (F2F), and NeuralTextures (NT).

where the videos have been processed with a *visually nearly lossless* compression [46]. This is more in line with the type of videos found on social media.

Table 1 shows that LipForensics achieves excellent generalisation to novel forgeries, surpassing on average most approaches by a large margin. It also outperforms the previous state-of-the-art method, Face X-ray, by 2.2% AUC, 94.9%  $\rightarrow$  97.1%. FaceSwap (to which learning-based methods tend not to generalise well, possibly due to its distinct visual artefacts) is the only manipulation for which Face X-ray performs better. Nonetheless, our approach manages to reach 90.1% AUC without imposing such strong *a priori* knowledge (*i.e.*, that there exists a blending boundary in FaceSwap frames).

We also notice that simply training the spatio-temporal network from scratch on mouth crops leads to surprisingly good results on the expression manipulation methods, *i.e.*, Face2Face and NeuralTextures. This is possibly because there are low-level temporal artefacts around the mouth that exist across the forgery types. More importantly, pretraining on lipreading improves performance across all methods, showing that the network now focuses on more transferable forgery evidence. Finally, freezing the feature extractor provides, on average, a further substantial increase in performance, suggesting that the extractor is prone to overfitting even after pretraining.

**Generalisation to other datasets.** To evaluate cross-dataset generalisation, we test a *single model* on DeeperForensics, Faceshifter, Celeb-DF-v2, and DFDC after training on FF++ (all four methods). In Table 2, we present the results for the same baselines as well as for three other methods: (1) **Multi-task** [43], which uses an autoencoder-like architecture similar to [13, 17]; (2) **DSP-FWA** [37]; and (3) **Two-branch** [41]. We train Multi-task ourselves on FaceForensics++, report the (video-level) results from the paper of Two-branch [41], and evaluate publicly-available,

Method	CDF	DFDC	FSh	DFo	Avg
Xception [46]	73.7	70.9	72.0	84.5	75.3
CNN-aug [59]	75.6	72.1	65.7	74.4	72.0
Patch-based [8]	69.6	65.6	57.8	81.8	68.7
Face X-ray [35]	79.5	65.5	92.8	86.8	81.2
CNN-GRU [47]	69.8	68.9	80.8	74.1	73.4
Multi-task [43]	75.7	68.1	66.0	77.7	71.9
DSP-FWA [37]	69.5	67.3	65.5	50.2	63.1
Two-branch [41]	76.7	—	—	—	—
Scratch	62.5	65.5	84.7	84.8	74.4
Ft whole	70.7	70.9	93.9	95.7	82.8
LipForensics (ours)	<b>82.4</b>	<b>73.5</b>	<b>97.1</b>	<b>97.6</b>	<b>87.7</b>

Table 2. **Cross-dataset generalisation.** Video-level AUC (%) on Celeb-DF-v2 (CDF), DeepFake Detection Challenge (DFDC), FaceShifter HQ (FSh), and DeeperForensics (DFo) when trained on FaceForensics++.

pretrained models for the remaining two methods.

LipForensics surpasses all methods on every dataset, with especially strong results on FaceShifter and DeeperForensics; this provides promise for its efficacy on future, improved forgeries. All methods obtain relatively low scores on DFDC (and to a lesser extent Celeb-DF), which we attribute to the domain shift caused by significantly different filming conditions. The gains from lipreading pretraining (+8.4% on average) and from freezing the feature extractor (+4.9%) are again apparent, providing additional strong evidence that high-level features related to mouth movements are key to generalisation.

### 4.3. Robustness to unseen perturbations

Given the ubiquity of image processing operations on social media, it is critical that deployed forgery detectors are not easily subverted by common perturbations. We investigate the robustness of the detectors by training on uncompressed FF++ and then testing on FF++ samples that were exposed to various unseen corruptions. We consider the following operations at five severity levels, as given in [30]: changes in saturation, changes in contrast, adding block-wise distortions, adding White Gaussian noise, blurring, pixelating, and applying video compression (H.264 codec). Figure 3 gives an example of each corruption at severity level 3. We remove the compression and pixelation training augmentations for Face X-ray and replace the Gaussian blurring augmentation for CNN-aug with median blurring, so that no test-time perturbation is seen during training.

In Figure 4, we show the effect of increasing the severity for each corruption. On Table 3, we give the average AUC scores across all severities for each corruption. It is evident that LipForensics is significantly more robust to most perturbations than other methods. For corruptions that affect the high frequency content of the frames (blur, pixelation, compression), it maintains high performance at all but the



Figure 3. **Corruption examples.** Examples of the corruptions considered in our robustness experiments at severity level 3; this set of corruptions was introduced in [30]. It consists of changes in saturation and contrast, block-wise distortions, white Gaussian noise, Gaussian blurring, pixelation, and video compression. More examples and information can be found in the supplementary material.

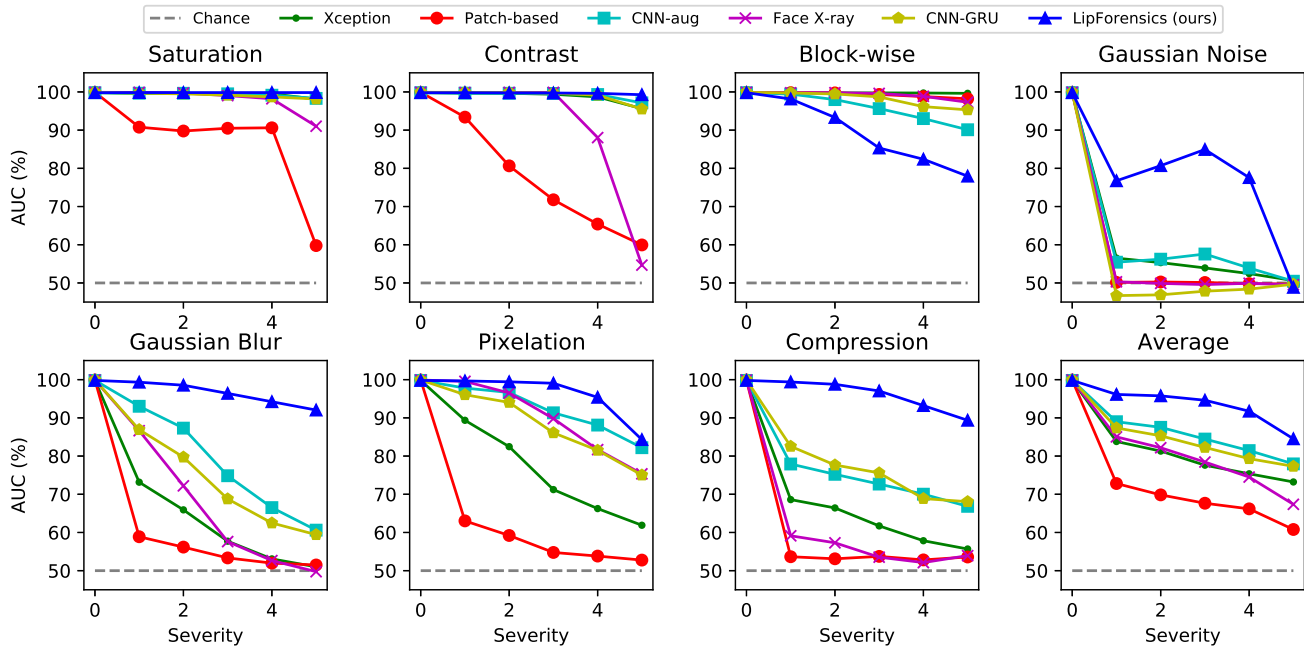


Figure 4. **Robustness to various unseen corruptions.** Video-level AUC scores as a function of the severity level for various corruptions. “Average” denotes the mean across all corruptions at each severity level. LipForensics is more robust than previous methods to all corruptions except block-wise distortions, which affect high-level semantic content.

most severe levels, while the other methods undergo significant deterioration in performance. Patch-based is particularly vulnerable to most corruptions; we attribute this to its reliance on a limited receptive field. Despite its good cross-manipulation performance, Face X-ray is adversely affected by most perturbations, especially compression, suggesting that the blending boundary is easily corruptible. It is also interesting that training with *median* blurring and *JPEG* compression augmentations (CNN-aug) is not an adequate remedy against *Gaussian* blurring and *video* compression. This agrees with CNN robustness literature [21, 55]. Finally, LipForensics is sensitive to block-wise distortions, which the other methods are relatively unaffected by. This result bolsters our hypothesis that other methods tend to focus on low-level cues, whereas LipForensics targets high-level inconsistencies, which block-wise distortions can destroy through occlusion. We also stress that such a corruption is

too conspicuous to be used adversarially and is also unlikely to be encountered in practice.

#### 4.4. Learning on compressed data

Next, we investigate in-distribution performance on FF++ at different levels of video compression. While we previously studied robustness to *unseen* corruptions, here we *train and test* a separate model for each of the video qualities provided by the dataset: (1) uncompressed videos (raw), (2) compressed videos at high quality (HQ), and (3) compressed videos at low quality (LQ) [46]. We also report, where applicable, results from the Two-branch paper [41].

Table 4 shows that while all models perform almost flawlessly on raw data, their efficacy varies when trained on compressed videos. Methods that employ deep networks (*e.g.*, Xception and CNN-aug) reach relatively high performance on the LQ dataset. In contrast, patch-based classi-

Method	Clean	Saturation	Contrast	Block	Noise	Blur	Pixel	Compress	Avg
Xception [46]	99.8	99.3	98.6	<b>99.7</b>	53.8	60.2	74.2	62.1	78.3
CNN-aug [59]	99.8	99.3	99.1	95.2	54.7	76.5	91.2	72.5	84.1
Patch-based [8]	99.9	84.3	74.2	99.2	50.0	54.4	56.7	53.4	67.5
Face X-ray [35]	99.8	97.6	88.5	99.1	49.8	63.8	88.6	55.2	77.5
CNN-GRU [47]	99.9	99.0	98.8	97.9	47.9	71.5	86.5	74.5	82.3
LipForensics (ours)	99.9	<b>99.9</b>	<b>99.6</b>	87.4	<b>73.8</b>	<b>96.1</b>	<b>95.6</b>	<b>95.6</b>	<b>92.5</b>

Table 3. **Average robustness to unseen corruptions.** Video-level AUC (%) on FF++ when videos are exposed to various corruptions, averaged over all severity levels. ‘‘Avg’’ denotes the mean across all corruptions (and all severity levels).

Method	Video-level acc (%)			Video-level AUC (%)		
	Raw	HQ	LQ	Raw	HQ	LQ
Xception [46]	99.0	97.0	89.0	99.8	99.3	92.0
CNN-aug [59]	98.7	96.9	81.9	99.8	99.1	86.9
Patch-based [8]	<b>99.3</b>	92.6	79.1	<b>99.9</b>	97.2	78.3
Two-branch [41]	—	—	—	—	99.1	91.1
Face X-ray [35]	99.1	78.4	34.2	99.8	97.8	77.3
CNN-GRU [47]	98.6	97.0	90.1	<b>99.9</b>	99.3	92.2
LipForensics (ours)	98.9	<b>98.8</b>	<b>94.2</b>	<b>99.9</b>	<b>99.7</b>	<b>98.1</b>

Table 4. **Learning on compressed data.** Performance on FF++ when trained and tested on uncompressed (Raw), slightly compressed (HQ), and heavily compressed (LQ) videos.

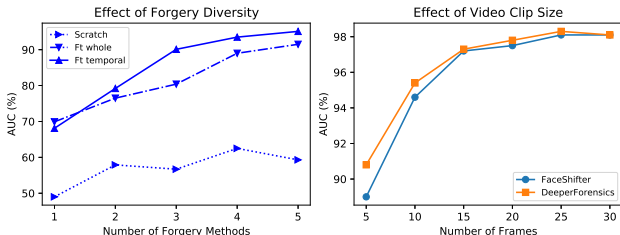


Figure 5. **Model analysis.** Effect of increasing forgery diversity on generalisation to FaceSwap (left; details in Section 4.5). Effect of video clip size on generalisation to FaceShifter and DeeperForensics when trained on FF++ (right).

fication struggles to effectively discriminate the data. We also notice that Face X-ray, albeit employing a very deep network, suffers the most under compression. This empirically validates our intuition that the blending artefacts are largely destroyed when the videos are compressed. On the other hand, LipForensics is substantially less affected by compressed data, outperforming all other methods, as it targets high-level spatio-temporal cues. We believe that the improvements over CNN-GRU and Two-branch, which also account for the temporal dimension, are due to the lipreading pretraining, the different architectures (convolutional versus recurrent), and the more frames per clip.

#### 4.5. Ablation study

**Effect of increasing forgery diversity.** With the development of new forgery generators, it will be possible to in-

Model	Finetune	FSh	DFo
ResNet-18	whole	59.3	75.7
ResNet-3D/2D	whole	67.1	74.6
ResNet-3D/2D+MS-TCN	whole	83.2	84.6
ResNet-3D/2D+MS-TCN (ours)	temporal	<b>87.5</b>	<b>90.4</b>

Table 5. **Effect of different components.** All models are pre-trained on LRW. We report video-level accuracy (%) scores on FaceShifter (FSh) and DeeperForensics (DFo) when finetuned on FaceForensics++. Last row corresponds to LipForensics.

crease the diversity of manipulation methods in the training set. We investigate how this affects generalisation. We treat FaceSwap as the unseen method, as it appears to be the most difficult for our detector to generalise to (see Table 1). We first train only on Deepfakes and then cumulatively add the following: NeuralTextures, FaceSwap, FaceShifter, and DeeperForensics (which share the same source videos).

In Figure 5, we see that LipForensics scales well with the forgery diversity, reaching 95.1% AUC on FaceSwap. Notably, finetuning the whole network still has inferior performance to freezing the feature extractor, even when training on five forgery methods. Without lipreading pretraining, the model does not achieve adequate generalisation.

**Influence of video clip size.** We next study the effect of video clip size on generalisation, when trained on FF++ (HQ). For the number of frames in each video clip, we consider the set  $\{5, 10, 15, 20, 25, 30\}$  and train a new model for each clip size. As Figure 5 shows, increasing the number of frames tends to improve performance. A size of 25 strikes a good balance between performance and computational/memory cost.

**Effect of different components.** We further study the importance of modelling high-level temporal inconsistencies. We pretrain on LRW a ResNet-18 (without a 3-D layer) followed by an MS-TCN, and then finetune only the ResNet-18 on FF++; thus, no temporal information can be exploited. We compare it with finetuning the ResNet-3D/2D, which can capture short-term, low-level temporal dynamics, and with the full spatio-temporal model that can model long-term, high-level irregularities. On Table 5, we see that

Model	# params	Pretrain	FSh	DFo
R(2+1)D-18	31.3M	none	63.6	65.4
R(2+1)D-18	31.3M	Kinetics	65.7	68.2
ip-CSN-152	32.2M	none	68.2	65.7
ip-CSN-152	32.2M	Kinetics	73.9	76.4
ip-CSN-152	32.2M	IG-65M	66.1	69.6
SE-ResNet50	43.8M	none	60.0	70.7
SE-ResNet50	43.8M	FR	64.3	68.9
ResNet+MS-TCN	36.0M	none	62.5	61.4
ResNet+MS-TCN	36.0M	LRW	83.2	84.6
ResNet+MS-TCN* (ours)	24.8M	LRW	<b>87.5</b>	<b>90.4</b>

Table 6. **Other pretraining datasets.** Effect of pretraining on Kinetics, IG-65M, and face recognition (FR) datasets. We give video-level accuracy scores (%) on FaceShifter and DeeperForensics when finetuned on FF++. “# params” refers to trainable parameters, and asterisk (\*) denotes freezing the feature extractor (corresponding to LipForensics).

modelling short-term dynamics improves generalisation to FaceShifter but not DeeperForensics. More importantly, *high-level* temporal information is crucial for generalisation, as indicated by the large improvements when the MS-TCN is added and also when the ResNet is kept frozen.

**Other pretraining datasets.** Does pretraining on other large-scale datasets work just as well as pretraining on lipreading? We note that most baselines considered thus far use ImageNet weights. Here, we further consider spatio-temporal models pretrained on Kinetics-400 [31] and IG-65M [22]. The former dataset comprises around 300,000 video clips spanning 400 action classes, and the latter is a massive-scale weakly-supervised dataset containing over 65 million social media clips with corresponding hashtags. We use the state-of-the-art R(2+1)D-18 [54] and ip-CSN-152 [53] models, which have a similar number of parameters as the ResNet-18+MS-TCN model we use. Furthermore, we compare with a Squeeze-and-Excitation [28] ResNet-50, pretrained using ArcFace loss [15] on an amalgamation of face recognition datasets<sup>3</sup>.

We finetune all models on FF++ and test on FaceShifter and DeeperForensics to evaluate cross-manipulation generalisation (see Table 6). Although other pretraining tasks generally improve performance over training from scratch, none of them match the generalisation achieved by pretraining on lipreading. This confirms its importance in inducing representations suitable for forgery detection.

**Occlusion sensitivity.** To visualise the spatial regions which the detectors rely on, we occlude different portions of the input by sliding a grey square or a square cuboid, and observe the change in the class probability, as proposed in

<sup>3</sup>[https://github.com/TreBlE\\_N/InsightFace\\_Pytorch](https://github.com/TreBlE_N/InsightFace_Pytorch)



Figure 6. **Occlusion sensitivity.** Visualisation of the regions which Xception (**top**) and LipForensics (**bottom**) rely on.

[62]. In Figure 6, we give results on randomly selected samples for LipForensics and Xception, both trained on FaceForensics HQ. As expected, LipForensics focuses predominantly on the mouth, which is not the case for Xception. More visualisations are in the supplementary material.

#### 4.6. Limitations

Despite its performance on various experiments, we acknowledge that LipForensics is not without limitations. It cannot be applied to isolated images, and it would probably not detect a fake video in which the mouth is occluded or left unaltered. However, arguably the most pernicious forgery content is videos where the mouth has been manipulated to modify speech, identity, or expression. It is also possible that there is a performance decline with limited mouth motion, although within most 25-frame clips, we noticed at least some movement, even in the absence of speech. Finally, it requires a large-scale labelled dataset for pretraining. Examples of failure cases are given in the supplementary material.

### 5. Conclusion

In this paper, we proposed a novel approach, dubbed *LipForensics*, for the detection of forged face videos. It targets inconsistencies in semantically high-level mouth movements by leveraging rich representations learned via lipreading. It achieves state-of-the-art generalisation to unseen forgery types while being significantly more robust than other methods to various common corruptions. Meeting both of these objectives is crucial for face forgery detection in real-life, and we believe that our work is an important step in the fight against fake videos.

**Acknowledgements.** We would like to thank Pingchuan Ma for his help with lipreading experiments. Alexandros Haliassos was financially supported by an Imperial President’s PhD Scholarship. All training, testing, and ablation studies have been conducted at Imperial College London.



## References

- [1] Deepfakes. <https://github.com/deepfakes/faceswap>. [Accessed: 2020-09-02]. 4
- [2] Faceswap. <https://github.com/MarekKowalski/FaceSwap>. [Accessed: 2020-09-03]. 4
- [3] Darius Afchar, Vincent Nozick, Junichi Yamagishi, and Isao Echizen. Mesonet: a compact facial video forgery detection network. In *2018 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–7. IEEE, 2018. 1, 2
- [4] Shruti Agarwal, Hany Farid, Ohad Fried, and Maneesh Agrawala. Detecting deep-fake videos from phoneme-viseme mismatches. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 660–661, 2020. 1, 2
- [5] Jawadul H Bappy, Amit K Roy-Chowdhury, Jason Bunk, Lakshmanan Nataraj, and BS Manjunath. Exploiting spatial structure for localizing manipulated image regions. In *Proceedings of the IEEE international conference on computer vision*, pages 4970–4979, 2017. 2
- [6] Belhassen Bayar and Matthew C Stamm. A deep learning approach to universal image manipulation detection using a new convolutional layer. In *Proceedings of the 4th ACM Workshop on Information Hiding and Multimedia Security*, pages 5–10, 2016. 2
- [7] Liron Bergman and Yedid Hoshen. Classification-based anomaly detection for general data. In *International Conference on Learning Representations (ICLR)*, 2020. 3
- [8] Lucy Chai, David Bau, Ser-Nam Lim, and Phillip Isola. What makes fake images detectable? understanding properties that generalize. *arXiv preprint arXiv:2008.10588*, 2020. 1, 2, 4, 5, 7
- [9] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014. 4
- [10] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017. 2
- [11] Komal Chugh, Parul Gupta, Abhinav Dhall, and Ramanathan Subramanian. Not made for each other-audio-visual dissonance-based deepfake detection and localization. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 439–447, 2020. 2
- [12] Joon Son Chung and Andrew Zisserman. Lip reading in the wild. In *Asian Conference on Computer Vision*, pages 87–103. Springer, 2016. 2, 4
- [13] Davide Cozzolino, Justus Thies, Andreas Rössler, Christian Riess, Matthias Nießner, and Luisa Verdoliva. Forensictransfer: Weakly-supervised domain adaptation for forgery detection. *arXiv preprint arXiv:1812.02510*, 2018. 1, 2, 3, 5
- [14] Hao Dang, Feng Liu, Joel Stehouwer, Xiaoming Liu, and Anil K Jain. On the detection of digital face manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5781–5790, 2020. 1
- [15] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019. 8
- [16] Brian Dolhansky, Joanna Bitton, Ben Pflaum, Jikuo Lu, Russ Howes, Menglin Wang, and Cristian Canton Ferrer. The deepfake detection challenge dataset. *arXiv preprint arXiv:2006.07397*, 2020. 1, 4
- [17] Mengnan Du, Shiva Pentylala, Yuening Li, and Xia Hu. Towards generalizable forgery detection with locality-aware autoencoder. *arXiv preprint arXiv:1909.05999*, 2019. 1, 2, 3, 5
- [18] Ricard Durall, Margret Keuper, and Janis Keuper. Watch your up-convolution: Cnn based generative deep neural networks are failing to reproduce spectral distributions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7890–7899, 2020. 2
- [19] Joel Frank, Thorsten Eisenhofer, Lea Schönherr, Asja Fischer, Dorothea Kolossa, and Thorsten Holz. Leveraging frequency analysis for deep fake image recognition. *arXiv preprint arXiv:2003.08685*, 2020. 2
- [20] T. Furlanello, Zachary Chase Lipton, Michael Tschannen, L. Itti, and Anima Anandkumar. Born again neural networks. In *International Conference on Machine Learning*, 2018. 4
- [21] Robert Geirhos, Carlos RM Temme, Jonas Rauber, Heiko H Schütt, Matthias Bethge, and Felix A Wichmann. Generalisation in humans and deep neural networks. In *Advances in neural information processing systems*, pages 7538–7550, 2018. 6
- [22] Deepti Ghadiyaram, Du Tran, and Dhruv Mahajan. Large-scale weakly-supervised pre-training for video action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12046–12055, 2019. 8
- [23] Izhak Golan and Ran El-Yaniv. Deep anomaly detection using geometric transformations. In *Advances in Neural Information Processing Systems*, pages 9758–9769, 2018. 3
- [24] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 1
- [25] David Güera and Edward J Delp. Deepfake video detection using recurrent neural networks. In *2018 15th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–6. IEEE, 2018. 1, 2
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3
- [27] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 2
- [28] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018. 8

- [29] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 4
- [30] Liming Jiang, Ren Li, Wayne Wu, Chen Qian, and Chen Change Loy. Deepforensics-1.0: A large-scale dataset for real-world face forgery detection. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2886–2895. IEEE, 2020. 1, 4, 5, 6
- [31] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 8
- [32] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 4
- [33] Pavel Korshunov and Sébastien Marcel. Speaker inconsistency detection in tampered video. In *2018 26th European Signal Processing Conference (EUSIPCO)*, pages 2375–2379. IEEE, 2018. 2
- [34] Lingzhi Li, Jianmin Bao, Hao Yang, Dong Chen, and Fang Wen. Advancing high fidelity identity swapping for forgery detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5074–5083, 2020. 1, 4
- [35] Lingzhi Li, Jianmin Bao, Ting Zhang, Hao Yang, Dong Chen, Fang Wen, and Baining Guo. Face x-ray for more general face forgery detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5001–5010, 2020. 1, 2, 3, 4, 5, 7
- [36] Yuezun Li, Ming-Ching Chang, and Siwei Lyu. In ictu oculi: Exposing ai created fake videos by detecting eye blinking. In *2018 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–7. IEEE, 2018. 1
- [37] Yuezun Li and Siwei Lyu. Exposing deepfake videos by detecting face warping artifacts. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2019. 2, 5
- [38] Yuezun Li, Xin Yang, Pu Sun, Honggang Qi, and Siwei Lyu. Celeb-df: A large-scale challenging dataset for deepfake forensics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3207–3216, 2020. 1, 2, 4
- [39] Pingchuan Ma, Brais Martinez, Stavros Petridis, and Maja Pantic. Towards practical lipreading with distilled and efficient models. *arXiv preprint arXiv:2007.06504*, 2020. 4
- [40] Brais Martinez, Pingchuan Ma, Stavros Petridis, and Maja Pantic. Lipreading using temporal convolutional networks. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6319–6323. IEEE, 2020. 4
- [41] Iacopo Masi, Aditya Killekar, Royston Marian Mascarenhas, Shenoy Pratik Gurudatt, and Wael AbdAlmageed. Two-branch recurrent network for isolating deepfakes in videos. *arXiv preprint arXiv:2008.03412*, 2020. 1, 2, 4, 5, 6, 7
- [42] Trisha Mittal, Uttaran Bhattacharya, Rohan Chandra, Aniket Bera, and Dinesh Manocha. Emotions don’t lie: An audio-visual deepfake detection method using affective cues. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2823–2832, 2020. 1, 2
- [43] H. H. Nguyen, F. Fang, J. Yamagishi, and I. Echizen. Multi-task learning for detecting and segmenting manipulated facial images and videos. In *2019 IEEE 10th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pages 1–8, 2019. 1, 2, 3, 4, 5
- [44] Yuval Nirkin, Yosi Keller, and Tal Hassner. Fsgan: Subject agnostic face swapping and reenactment. In *Proceedings of the IEEE international conference on computer vision*, pages 7184–7193, 2019. 1
- [45] Yuyang Qian, Guojun Yin, Lu Sheng, Zixuan Chen, and Jing Shao. Thinking in frequency: Face forgery detection by mining frequency-aware clues. In *European Conference on Computer Vision*, pages 86–103. Springer, 2020. 1, 2
- [46] Andreas Rossler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. Faceforensics++: Learning to detect manipulated facial images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1–11, 2019. 1, 2, 4, 5, 6, 7
- [47] Ekraam Sabir, Jiaxin Cheng, Ayush Jaiswal, Wael AbdAlmageed, Iacopo Masi, and Prem Natarajan. Recurrent convolutional strategies for face manipulation detection in videos. *Interfaces (GUI)*, 3(1), 2019. 2, 4, 5, 7
- [48] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5693–5703, 2019. 4
- [49] Michail Tarasiou and Stefanos Zafeiriou. Extracting deep local features to detect manipulated images of human faces. In *2020 IEEE International Conference on Image Processing (ICIP)*, pages 1821–1825. IEEE, 2020. 2
- [50] Justus Thies, Mohamed Elgharib, Ayush Tewari, Christian Theobalt, and Matthias Nießner. Neural voice puppetry: Audio-driven facial reenactment. In *European Conference on Computer Vision*, pages 716–731. Springer, 2020. 1
- [51] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019. 1, 4
- [52] Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Face2face: Real-time face capture and reenactment of rgb videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2387–2395, 2016. 1, 4
- [53] Du Tran, Heng Wang, Lorenzo Torresani, and Matt Feiszli. Video classification with channel-separated convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5552–5561, 2019. 8
- [54] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6450–6459, 2018. 8

- [55] Igor Vasiljevic, Ayan Chakrabarti, and Gregory Shakhnarovich. Examining the impact of blur on recognition by convolutional networks. *arXiv preprint arXiv:1611.05760*, 2016. [6](#)
- [56] Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. End-to-end speech-driven realistic facial animation with temporal gans. In *CVPR Workshops*, pages 37–40, 2019. [1](#)
- [57] Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Realistic speech-driven facial animation with gans. *International Journal of Computer Vision*, pages 1–16, 2019. [1](#)
- [58] Run Wang, Felix Juefei-Xu, Lei Ma, Xiaofei Xie, Yihao Huang, Jian Wang, and Yang Liu. Fakespotter: A simple yet robust baseline for spotting ai-synthesized fake faces. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2020. [2](#)
- [59] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. Cnn-generated images are surprisingly easy to spot... for now. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 7, 2020. [1](#), [2](#), [3](#), [4](#), [5](#), [7](#)
- [60] Xinsheng Xuan, Bo Peng, Wei Wang, and Jing Dong. On the generalization of gan image forensics. In *Chinese Conference on Biometric Recognition*, pages 134–141. Springer, 2019. [3](#)
- [61] Chen-Zhao Yang, Jun Ma, Shi-Lin Wang, and Alan Wee-Chung Liew. Preventing deepfake attacks on speaker authentication by dynamic lip movement analysis. *IEEE Transactions on Information Forensics and Security*, 2020. [2](#)
- [62] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014. [8](#)
- [63] Xu Zhang, Svebor Karaman, and Shih-Fu Chang. Detecting and simulating artifacts in gan fake images. In *2019 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–6. IEEE, 2019. [2](#)
- [64] Peng Zhou, Xintong Han, Vlad I Morariu, and Larry S Davis. Two-stream neural networks for tampered face detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1831–1839. IEEE, 2017. [1](#)