

This CVPR 2021 workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Adversarial Threats to DeepFake Detection: A Practical Perspective

Paarth Neekhara UC San Diego Brian Dolhansky Facebook AI Joanna Bitton Facebook AI Cristian Canton Ferrer Facebook AI

Abstract

Facially manipulated images and videos or DeepFakes can be used maliciously to fuel misinformation or defame individuals. Therefore, detecting DeepFakes is crucial to increase the credibility of social media platforms and other media sharing web sites. State-of-the art DeepFake detection techniques rely on neural network based classification models which are known to be vulnerable to adversarial examples. In this work, we study the vulnerabilities of state-of-the-art DeepFake detection methods from a practical stand point. We perform adversarial attacks on DeepFake detectors in a black box setting where the adversary does not have complete knowledge of the classification models. We study the extent to which adversarial perturbations transfer across different models and propose techniques to improve the transferability of adversarial examples. We also create more accessible attacks using Universal Adversarial Perturbations which pose a very feasible attack scenario since they can be easily shared amongst attackers. We perform our evaluations on the winning entries of the DeepFake Detection Challenge (DFDC) and demonstrate that they can be easily bypassed in a practical attack scenario by designing transferable and accessible adversarial attacks.¹

1. Introduction

DeepFakes are artificial videos that contain realistically swapped faces mostly created with off-the-shelf neural network based methods. While DeepFakes are sometimes used for humorous or entertainment purposes, these videos are an emerging threat, especially within the realms of politics and misinformation [43]. DeepFakes are especially convincing and have caused harm by making it appear that a prominent person said or did something that they never said or did. Accordingly, effort has been devoted towards training DeepFake classifiers and detectors, which attempt to determine if a video contains a fake face via a variety of methods [42]. The state-of-the-art DeepFake detection methods rely on Convolutional Neural Networks (CNNs) to classify a given video as *Real* or *Fake* [11]. The best performing methods model the DeepFake detection problem as a per-frame classification problem. While such methods achieve promising results in terms of detection accuracy, they are vulnerable to adversarial examples [40] and can be evaded by adding a carefully crafted perturbation to each frame of a given input video [5, 13, 24]. Since DeepFakes have the potential to be very damaging, attacks designed to evade DeepFake detectors can cause outsized harm when compared to other attack scenarios. In addition, as DeepFakes are rare compared to the set of all videos, detection of DeepFakes is already an extremely difficult problem.

While adversarial examples pose a threat to DeepFake detectors, designing such examples usually requires complete access to the victim detector model architecture and parameters. In a practical threat scenario, the model weights can be kept secret to prevent such white-box attacks. While past works have also proposed black-box attacks to Deep-Fake video detectors [24], they require a large number of queries and access to the model prediction scores for each frame in the video that they aim to misclassify. Such an attack can easily be thwarted by limiting query access and not providing the raw detection scores to the user.

Adversarial examples can pose a practical threat to DeepFake detection if they are transferable across different models. Past works have shown that adversarial examples designed to fool a particular network can also fool other networks (with the same or different architecture) trained for the same task [14, 39]. By exploiting this property, an adversary can design attacks on an open source DeepFake detection model and potentially fool a DeepFake detection system in production. However, in a real-word scenario, different detection mechanisms employ different input preprocessing steps which can nullify many adversarial attacks. Additionally, DeepFake detectors also use face detection techniques prior to image classification CNNs which can differ across various detection methods. This makes it challenging to craft perturbations that are transferable across different detection methods.

¹Video Examples: https://deepfakeattacks.github.io/

The goal of our work is to study the practical threats posed by adversarial examples to the current state-of-the art DeepFake detection systems. To this end, we first study the commonalities between different DeepFake detection methods and gain insight by interpreting the model decisions using gradient-based saliency maps. We then study the vulnerability of these detection methods to adversarial examples and the extent to which adversarial examples transfer across different detection methods. Next, we propose attack methods that significantly improve the transferability of adversarial examples by designing perturbations that are robust to the differences between the various detection methods. Finally, we design more accessible adversarial attacks by creating transferable universal adversarial perturbations that can be universally added across all frames of all videos to reliably fool a number of DeepFake detection methods.

2. Background

2.1. DeepFakes and DFDC

DeepFakes are a genre of synthetic videos in which a subject's face is swapped with a target face to simulate the target subject in a certain scenario and create convincing footage of events that never occurred [34, 24]. Recent video manipulation methods operate end-to-end on a source video and target face and require minimal human expertise to generate fake videos in real-time [16, 44, 26]. For expediting research on DeepFake detection, there has been effort in curating datasets [11, 34] of real and fake videos using Deep-Fake synthesis techniques.

To the best of our knowledge, the recently developed DeepFake Detection Challenge (DFDC) dataset [12, 11] is the largest collection of such real and fake videos, consisting of over 1 million training clips of face swaps produced with a variety of methods. For synthesizing the fake videos in the DFDC dataset, 8 different video manipulation techniques were used, many of which are CNNbased techniques. These methods include the traditional DeepFake auto-encoder architecture, a non-learned morphable mask face swap algorithm, and several Generative Adversarial Networks (GAN) techniques like Neural Talking Heads [44], FSGAN [26] and StyleGAN [16]. In conjunction with the dataset, a corresponding competition² was launched in which competitors were encouraged to submit models trained for DeepFake detection on the training set. These models were then ranked on a hidden, held-out test set, and the winning competitors released their architectures and training strategies publicly.

2.2. DeepFake Detection

Recent state-of-the-art methods for detecting manipulated facial content in videos rely on Convolutional Neural Networks [1, 2, 19, 32, 34] to distinguish AI-generated fake videos from real videos. These methods model the DeepFake identification problem as a per-frame classification problem. They employ typical image classification networks that either operate on the entire frame or on a cropped portion of the frame that has domain specific information. For example, state-of-the-art DeepFake classification systems [9, 15, 35] consist of a face-tracking method, following by the cropped face passed on to a CNN-based classifier for classification as Real or Fake [1, 8]. The final label of a video is usually the aggregation of the labels for some candidate frames of the video.

While it seems intuitive that exploiting temporal dependencies using sequence models should improve a detector's ability to spot manipulated videos, the insights from the results of the DFDC challenge [11, 12] show that the best performing models operate on a frame level. In fact, the winning team [35] of the DFDC challenge explicitly noted that ideas besides the frame-by-frame detector did not improve their performance on the public leader-board. One reason we anticipate for this is that the recent DeepFake generation techniques have improved temporal consistencies in the videos; however creating plausible face-swaps in images is still challenging due to the artifacts introduced by upsampling methods in autoencoders.

In our work, we focus on the top three winning entries [9, 15, 35] of the DFDC challenge. In order to understand their vulnerabilities, we first studied the commonalities across these detection methods. Our objective was to gain insight into what the detector is looking at when it makes a decision about a video being Real or Fake. This is typically done by obtaining the gradient of the score of the predicted class with respect to the input image and plotting the magnitude of these gradients as a heat-map. Back-propagating gradients naively does not result in interpretable visualizations. This is because we only care about pixels which activate a neuron rather than suppress it (suppression is indicated by negative gradients) [38]. Therefore, we use guided backpropagation which defines custom gradient estimates for activation functions like ReLU and suppresses negative gradients during the backward pass [38]. We then standardize the gradient obtained with respect to the input and overlay the heat-map on the frame to visualize the areas of the image that trigger the network's output. Figure 1 shows some examples of the saliency maps obtained while analyzing two different detectors on DeepFake videos.

!

Our initial observations on these saliency maps suggested that different CNN-based detection methods attend to similar aspects of the input frame for predicting the label. These aspects include the edges of the face, the eyes, lips, teeth etc. These similarities across different detection

²https://www.kaggle.com/c/deepfake-detection-challenge



Figure 1. Gradient saliency maps obtained on DeepFake videos using guideded-backpropogation on a CNN-based detector [35]. The highlighted areas indicate the image regions that strongly influence the detector's predictions.

methods indicate that adversarially modifying such aspects of the image could potentially fool multiple detection methods. We validate this hypothesis in our work by studying the transferability of adversarial examples (Section 4) across different detection methods and proposing techniques (Section 3.4) that improve the transferability.

2.3. Prior Work on Fooling DeepFake Detectors

Neural network-based classification systems have been known to be vulnerable to adversarial examples. Adversarial examples are intentionally designed inputs to a machine learning (ML) model that cause the model to make a mistake [40]. Prior work has shown that gradient-based attacks can effectively fool neural networks [6, 14, 22, 29, 30, 36, 37] with a minimal amount of perturbation added to the original input.

Recently, gradient-based adversarial attacks have also been applied on CNN-based DeepFake detection systems to expose their vulnerabilities to adversarial examples [5, 13, 24]. While some of this past work [5, 13] focuses on attacks on image classification models, the authors of [24] study the vulnerability of video DeepFake detection methods which follow the same detection pipeline as the methods studied in our work. While this past work demonstrates that adversarial examples can fool video DeepFake detectors, designing such adversarial videos requires complete access to the victim model architecture and parameters (white-box attack). This assumption makes the threat very limited in a realworld scenario since the model architecture and parameters can be kept hidden by the service provider. The black-box attacks proposed in [24] require querying the victim model multiple times and access to the raw scores given by the classifier for each frame the attacker wishes to misclassify. The number of queries and access to raw classifier scores can be restricted to thwart the adversary thereby preventing this black-box attack.

Adversarial examples pose a practical threat to Deep-Fake detection if they are transferable across different de-



Figure 2. A typical DeepFake detection pipeline. A face-tracking model crops the face from all video frames, which is resized and normalized appropriately to be fed as input to a CNN classification model.

tection methods. That is, if adversarial videos designed to fool some open source DeepFake detection method can also reliably fool other unseen CNN-based detection methods, it poses a real security threat to deploying CNN-based detectors in production. Several past works have studied this *transferability* property of adversarial examples where an attacker first generates an adversarial perturbation on a (white-box) surrogate source model, and then transfers it to the unknown target network [7, 14, 21, 27, 28, 39, 46].

In our work, we study the transferability of adversarial examples across various DeepFake detection methods. We find that differences in 1) input-preprocessing steps and 2) face detection methods across DeepFake detectors hamper the cross-model transferability of adversarial examples. We propose techniques to overcome these challenges in Section 3.4, and further propose more accessible attacks inspired from universal adversarial perturbations [4, 20, 23, 25, 33]. Universal adversarial perturbations pose a more practical threat to DeepFake detection since they can be easily shared amongst attackers and require no technical expertise in adversarial machine learning.

3. Methodology

3.1. Threat Model

Given a video (*Real* or *Fake*), our task is to adversarially modify the video such that the label predicted by a victim DeepFake detection method is incorrect. That is, we want to modify the videos such that the *Fake* videos are classified as *Real* and vice-versa. Misclassifying a *Fake* video as *Real* can be used by the adversary to propagate false information. Misclassifying a *Real* video as *Fake* can be used by the adversary to cover up an event that did actually happen.

The goal is to also ensure imperceptibility of the adversarial perturbation. In the image domain, L_p norms are commonly used to quantify the amount of perturbation added to create an adversarial input. The authors of [14] recommend constraining the maximum distortion of any individual pixel in the image using the L_{∞} metric. In the attacks discussed in this paper, the adversarial perturbation is added to each frame of the input video to create the adversarially modified video. Following past work on fooling DeepFake detectors [24], we use the L_{∞} metric to con-

strain the the amount of distortion added to each frame.

Notation: We follow the notation previously used in [6, 31, 24]; we define F to be the full neural network (classifier) including the softmax function, Z(x) = z to be the output of all layers except the softmax (that is z are the logits). That is:

$$F(x) = softmax(Z(x)) = y$$

The classifier assigns the label $C(x) = \arg \max_i (F(x)_i)$ to input frame x.

Problem Formulation: Mathematically, for each video frame x, we aim to find an adversarial frame x_{adv} such that:

$$C(x_{adv}) = y$$
 and $||x_{adv} - x_0||_{\infty} < \epsilon$

where y is the target label. In our case the target label is *Real* for *Fake* videos and *Fake* for *Real* videos. In the upcoming sections, we study this attack goal in various attacker knowledge settings and constraints.

3.2. Victim Models

The victim models we consider in our work model Deep-Fake detection as a per-frame classification problem. These models further decompose the frame classification problem into the following two steps:

- 1. A face tracking model detects the bounding box of the face in each frame.
- 2. The cropped face is then pre-processed using some input transformations (e.g. resizing, center-cropping and normalization) and fed as input to a CNN classification model that scores the frame as *Real* or *Fake*.

Finally, the scores of all or a subset of the frames are aggregated to obtain the final label of the video. The above detection pipeline has been used by the top 5 winning entries of the DFDC challenge. The CNN architectures and data-augmentation procedures vary across different methods. Table 1 details various DeepFake detection methods considered in our work with their respective face detection methods and CNN architectures used for classification.

3.3. White-box attacks

In this setting, we assume that the attacker has complete access to the detection model, including the face extraction pipeline and the architecture and parameters of the classification model. To construct adversarial examples using the attack pipeline described above, we use the iterative gradient sign method [17] to optimize the following objective:

Minimize
$$loss(x')$$
 where
 $loss(x') = max(Z(x')_o - Z(x')_y, 0)$
(1)

Here, $Z(x)_y$ is the final score for target label y and $Z(x)_o$ is the score of the original label o before the softmax operation in the classifier C. The loss function we

use is recommended by [6] because it is empirically found to generate less distorted adversarial samples and is robust against defensive distillation. We use the iterative gradient sign method to optimize the above loss function while constraining the magnitude of the perturbation as follows:

$$x_i = x_{i-1} - \operatorname{clip}_{\epsilon}(\alpha \cdot \operatorname{sign}(\nabla \operatorname{loss}(x_{i-1})))$$
(2)

We continue gradient descent iterations until success or until a given number number of maximum iterations, whichever occurs earlier. We solve the optimization problem for each frame of the given video and combine all the adversarial frames together to generate the adversarial video. In our experiments, we demonstrate that we are able to successfully fool all the detection methods studied in our work in the white-box attack setting using the above attack. However, the transferability of adversarial examples generated using this attack across different methods is limited. In the next section we propose techniques to overcome this challenge.

3.4. Black-box: Transfer attacks

Past works (Section 2.3) have studied that adversarial inputs can transfer across different models. That is, an adversarial input that was designed to fool a particular victim model can possibly fool other models that were trained for the same task. This is because different models learn similar decision boundaries and therefore have similar vulnerabilities. However, for DeepFake detectors, the goal of making transferable adversarial videos is more challenging due to multiple steps involved in the DeepFake detection pipeline and the differences in these steps across various methods.

- Different face detection methods result in different face-crops.
- Different data-augmentation procedures during training result in different levels of robustness to adversarial examples.
- Different input pre-processing pipelines, such as image resizing, cropping and normalization parameters, vary across different detection methods.

Therefore, to craft *transferable* adversarial videos, it is important to ensure robustness to such differences across various methods. To accomplish this, we craft adversarial examples that are robust over a given distribution of input transformations [3]. Given a distribution of input transformations T, input image x, and target class y, our objective is as follows:

$$x_{adv} = argmax_x \mathbb{E}_{t \sim T}[F(t(x))_y] \text{ s.t. } ||x - x_0||_{\infty} < \epsilon$$

That is, we want to maximize the expected probability of target class y over the distribution of input transforms T. To

solve the above problem, we update the loss function given in Equation 1 to be an expectation over input transforms T as follows:

$$loss(x) = \mathbb{E}_{t \sim T}[max(Z(t(x))_o - Z(t(x))_y, 0)]$$

Following the law of large numbers, we estimate the above loss functions for n samples as:

$$loss(x) = \frac{1}{n} \sum_{t_i \sim T} [max(Z(t_i(x))_o - Z(t_i(x))_y, 0)] \quad (3)$$

Since the above loss function is a sum of differentiable functions, it is tractable to compute the gradient of the loss w.r.t. to the input x. We minimize this loss using the iterative gradient sign method given by Equation 2. We iterate until a maximum number of iterations is reached or until the attack is successful under the sampled set of transformation functions, whichever happens first.

Next, we describe the class of input transformation functions we consider for the distribution T:

- Translation: We pad the image on all four sides by zeros and shift the pixels horizontally and vertically by a given amount. This transform ensures robustness to different face-detection and cropping pipelines across various methods. Let t_x be the transform in the x axis and t_y be the transform in the y axis, then $t(x) = x'_{H,W,C}$ s.t $x'[i, j, c] = x[i + t_x, j + t_y, c]$
- Downsizing and Upsizing: The image is first downsized by a factor r and then up-sampled by the same factor using bilinear re-sampling.
- Gaussian Noise Addition: Addition of Gaussian noise sampled from Θ ~ N(0, σ) to the input image. This transform is given by t(x) = x + Θ

The details of the hyper-parameter search distribution used for these transforms can be found in Section 4.3. Empirically, we find that ensuring robustness of adversarial examples significantly improves attack transferability across various detection methods.

3.5. Universal attacks

While the transferability of adversarial perturbations poses a practical threat to DeepFake detectors in production, creating an adversarial video requires significant technical expertise in adversarial machine learning - the attacker needs to solve an optimization problem for each frame of the video to fool the detector.

To ease the process of fooling DeepFake detectors, we aim to design more accessible adversarial attacks that can be easily shared amongst attackers. Past works [22, 4, 25] have shown the existence of universal adversarial perturbations that can fool classification models in various input domains. We aim to find a single universal adversarial perturbation which when added across all frames of any video, will cause the victim DeepFake Detector to classify the video to a target label.

That is, we aim to find a targeted universal perturbation δ such that:

$$C(x+\delta) = y \quad s.t \quad ||\delta||_{\infty} < \epsilon$$

for "most" x in our dataset (4)

where y is the target class. We train separate perturbations for Real and Fake target labels. In order to ensure robustness to differences across detection methods, we incorporate the transformation functions described in Section 3.4. We train the universal adversarial perturbation on a dataset of videos that are labelled opposite from our target label. On this dataset of videos, we aim to maximize the log-likelihood of predicting our target label y. Additionally to ensure the imperceptibility of the adversarial perturbation we penalize the L_2 distortion of the perturbation by adding a regularization term in our objective. Thus, our final objective to train the a universal perturbation for a target label y is as follows:

$$\begin{aligned} \text{Minimize} \sum_{x \text{ in } D} \mathbb{E}_{t \sim T}[L(F(t(x+\delta)), y)] + c ||\delta||_2 \\ \text{such that} \quad ||\delta||_{\infty} < \epsilon \end{aligned} \tag{5}$$

Here, L is the cross-entropy loss between the predictions and our target label, c is a hyper-parameter to control the regularization loss and x is an input frame of a video from our dataset D. Similar to Equation 3, we estimate the above expectation using n samples as follows:

$$\mathbb{E}_{t\sim T}[L(F(t(x+\delta)), y)] = \frac{1}{n} \sum_{t_i \sim T} [L(F(t_i(x+\delta)), y)]$$
(6)

To ensure the constrain $||\delta||_{\infty} < \epsilon$, we express δ as follows:

$$\delta = \epsilon \cdot tanh(p)$$

where p is a trainable unconstrained parameter having the same dimensions as δ . We fix the size of the perturbation vector p to be $3 \times 256 \times 256$ in our experiments, but resize the perturbation using bilinear interpolation to match the size of our input x. We iteratively optimize the objective given by Equation 5 using gradient descent. In our experiments, we find that targeting certain DeepFake detectors not only results in input-agnostic universal perturbations but also model-agnostic universal perturbations.

4. Experiments

4.1. Experimental Setup

DeepFake detectors: In our work, we consider the Deep-Fake detection methods proposed by the top three win-

Model	Team Name	Classifier	Face detection	AUC
EN-B7 Selim [35]	Selim	EfficientNet B7 [41]	MTCNN [45]	0.717
XN WM [15]	Team WM	XceptionNet [8]	RetinaFace [10]	0.724
EN-B3 WM [15]	Team WM	EfficientNet B3 [41]	RetinaFace [10]	0.724
EN-B7 NLab [9]	NTech Lab	EfficientNet B7 [41]	DSFD [18]	0.717

Table 1. Different DeepFake detection systems studied in our work with their respective classification models, face detection models and detection AUC scores on the DFDC test set.

ning entries of DFDC [12]. All of these methods follow the DeepFake detection pipeline described in Section 3.2. However, these methods use different CNN model architectures for classification and face-tracking. Table 1 lists the DeepFake detection methods studied in this work along with their respective CNN architectures used for classification and face detection. In our experiments we use the terms *victim model* and *test model* and define them as:

- *Victim model:* The detection model that the attack/adversarial perturbation is trained on, in the complete-knowledge (white-box) attack scenario.
- *Test model:* The model on which we evaluate the attack
 Can be the same as the victim model (white-box) or an unseen detection model (black-box).

Datasets: We craft adversarial videos for the first 100 Fake and 100 Real videos in the public DFDC validation set [12]. These videos contain a total of 30,300 frames. The videos are recorded in various lighting and background conditions and include people with different skin-tones.

Evaluation Metrics: After performing our attacks we combine the adversarial frames to create the adversarial video. We report the following metrics for evaluating our attacks:

Success Rate (SR): The percentage of videos for which we are able to successfully flip the original correct label predicted by a given detection method. Note that we do not take into account the videos that are originally mis-predicted by the classifier.

Mean distortion (L_{∞}) : The average L_{∞} distortion between the adversarial and original frames. The pixel values are scaled in the range [0,1], so changing a pixel from full-on to full-off in a grayscale image would result in L_{∞} distortion of 1 (not 255).

4.2. White-box attacks

In this section, we discuss the evaluation of the iterative gradient based white-box attack described in Section 3.3. Following past work [24], we set the max allowed L_{∞} norm ϵ as 16/255 and continue the attack iterations until the predicted score of our target label is greater than 0.99.

As shown in Table 3, for a given victim model, we are able to achieve 100% success rate for the same test model.



Figure 3. Examples of benign and adversarial video frames generated using our simple white-box attack targeting different Deep-Fake detection methods. It can be seen that the original prediction and label of the video can be successfully flipped by adding an imperceptible amount of adversarial perturbation in a white-box attack setting.

EfficientNet-B7 by NTech Lab requires the highest amount of adversarial perturbation under the L_{∞} metric as compared to other methods in this study. We also evaluate the extent to which these perturbations transfer across different methods. We find that perturbations trained to fool EfficientNet-B7 by Team NTech Lab result in the most transferable attacks as indicated by the higher success rates on other test models. This suggests that *EN-B7 NLab* is relatively more robust to adversarial perturbations in comparison to the other models used in this study (also indicated by higher L_{∞} perturbation required to fool *EN-B7 NLab*).

Figure 3 shows examples of adversarial faces generated to flip the original label of a given frame while targeting different victim models. As indicated by the low L_{∞} norm of the perturbation, the amount of added perturbation is very imperceptible indicating that DeepFake detectors are extremely vulnerable to adversarial examples and can be easily fooled. Video examples are linked in the footnote on the first page.

4.3. Transfer attacks

To improve the transferability of adversarial examples across different methods, we perform our transfer attack described in Section 3.4 and evaluate the adversarial videos against unseen detection methods in a black-box setting. The hyper-parameters of the transformation functions used for the attack have been provided in Table 2. All other attack hyper-parameters are kept the same as our simple white-box attack.

As indicated by the results in Table 4, we are able to significantly improve the transferability of adversarial perturbations across different detection methods as compared to our simple white-box attack. The adversarial perturba-

Transform	Hyper-parameter search distribution
Translation	$ d_x \sim \mathcal{U}[-20, 20], d_y \sim \mathcal{U}[-20, 20]$
Gaussian Noise	$\sigma \sim \mathcal{U}[0.05, 0.07]$
Down-sizing & Up-sizing	Scaling factor $r \sim \mathcal{U}[2,5]$

Table 2. Search distribution of hyper-parameters of different transformations used for our transfer attack. During training, we sample three functions from each of the transforms to estimate the gradient of our expectation over transforms.



Figure 4. Examples of benign and adversarial video frames generated using our transfer attack targeting different DeepFake detection methods.

tions are most transferable across models with the same architecture. For example, we are able to achieve high crosstransferability between *EN-B7 Selim* vs *EN-B7 NLab*. Similar to our observation in the previous section, attacking *EN-B7 NLab* results in the most transferable adversarial attacks - we are able to achieve at least 72% success rate across all other detection methods when attacking *EN-B7 NLab*.

In order to ensure the robustness of the adversarial perturbation to input transformations, a relatively higher amount of perturbation is required as compared to our white-box attack (Mean L_{∞} distortion 0.0135 vs 0.0077). Figure 4 shows examples of adversarial faces generated to flip the original label of a given frame while targeting different victim models using our transfer attack. While a higher amount of perturbation is required as compared to our simple white-box attack, visually the perturbation is still fairly imperceptible.

4.4. Universal attacks

To create more accessible attacks, we train a universal adversarial perturbation using the procedure described in Section 3.5. We set the L_2 regularization term c = 0.01 and use the Adam optimizer with a learning rate of 0.001. For our initial experiments, we set the the L_{∞} threshold $\epsilon = 40/255$ for all victim models. Since the goal of find-

	Test Models				
Victim Model	$ L_{\infty} $	EN-B7 Selim	EN-B7 NLab	XN WM	EN-B3 WM
EN-B7 Selim	0.007	100.0%	59.5%	57.0%	38.5%
EN-B7 NLab	0.013	94.0%	100.0%	66.5%	49.5%
XN WM	0.006	13.0%	12.5%	100.0%	12.0%
EN-B3 WM	0.005	21.0%	15.5%	22.0%	100.0%

Table 3. Attack success rates of the white-box attacks (Section 3.3) on different victim models and their transferability to unseen detectors (test models).

	Test Models				
Victim Model	$ L_{\infty} $	EN-B7 Selim	EN-B7 NLab	XN WM	EN-B3 WM
EN-B7 Selim	0.010	100.0%	89.0%	72.5%	62.0%
EN-B7 NLab	0.018	99.0%	100.0%	72.0%	76.5%
XN WM	0.018	49.0%	33.5%	100.0%	46.0%
EN-B3 WM	0.008	46.5%	35.0%	47.5%	100.0%

Table 4. Attack success rates of the transfer attacks (Section 3.4) on different victim models and their transferability to unseen detectors (test models).

	Test Models				
Victim Model	$ L_{\infty} $	EN-B7 Selim	EN-B7 NLab	XN WM	EN-B3 WM
EN-B7 Selim	0.156	100.0%	94.5%	65.0%	68.75%
EN-B7 NLab	0.156	94.5%	100.0%	75.0%	81.50%
XN WM	0.156	77.5%	61.0%	100.0%	20.0%
EN-B3 WM	0.156	66.5%	50.5%	60.0%	100.0%

Table 5. Attack success rates of the universal attacks (Section 3.5) on different victim models and their transferability to unseen detectors (test models).

ing a single input-agnostic perturbation is more challenging than finding one perturbation per video frame, a higher amount of distortion is required for a successful attack as compared to the per-frame attacks described earlier. We train the universal perturbation on a dataset of 100 videos from the DFDC train set which are separate from our evaluation dataset. We train the perturbation using a batch size of 8 for 10,000 iterations.

We target one victim model at a time and test the transferability of the universal perturbation on seen and unseen detectors. Table 5 presents the results of performing the universal attack on different victim models at $\epsilon = 40/255 =$ 0.156. We are able to achieve 100% attack success rate on the same test model as the victim model using a single perturbation across all frames and videos of the same label. Also, the universal perturbation is transferable to a significant extent across different models which poses an extremely practical threat to DeepFake detectors in production. Attacking *EN-B7 NLab* results in the most transferable perturbations where we are able to achieve at least a 75% success rate across all unseen detectors.

Visually, the universal perturbations at $\epsilon = 0.156$ are more perceptible than our per-frame attacks discussed in



Figure 5. Visualization of universal adversarial perturbations trained on different DeepFake detection models.

the sections above. Figure 5 shows examples of universal adversarial perturbations trained on different DeepFake detectors and the resulting adversarial images obtained after adding the perturbation to the face-crop of the benign frame.

We perform an additional experiment to study the effectiveness of universal adversarial perturbations at different magnitudes of added perturbations. We choose *EN-B7 NLab* as the victim model and perform our universal attack at different values of ϵ . The attack success rates across different models are shown in Figure 6. Figure 7 shows what a perturbed image looks like at different values of ϵ . At $\epsilon < 0.1$, the perturbation is fairly imperceptible but can still achieve high success rates on various test models.

5. Conclusion

We design transferable black-box attacks which pose a practical threat to the security of DeepFake detection. Through our design of universal adversarial perturbations, we demonstrate the ease of accessibility of such attacks since the same perturbation can be added to all the frames of any video in order to bypass multiple detection systems. By



Figure 6. Attack success rate on unseen detectors for a universal perturbation trained on the *EN-B7 NLab* detector at different levels of the L_{∞} norm of the perturbation.



Figure 7. Visualization of the perturbed images using different magnitudes (ϵ) of universal adversarial perturbations trained on *EN-B7 NLab*.

bypassing recently proposed state-of-the-art DeepFake detection networks with our proposed attack algorithms, this work emphasizes the need for stronger detection methods that are robust to adversaries.

References

- Darius Afchar, Vincent Nozick, Junichi Yamagishi, and Isao Echizen. MesoNet: a compact facial video forgery detection network. In *Workshop on Information Forensics and Security* (*WIFS*), 2018. 2
- [2] Irene Amerini, Leonardo Galteri, Roberto Caldelli, and Alberto Del Bimbo. Deepfake video detection through optical flow based CNN. In *International Conference on Computer Vision (ICCV) Workshops*, 2019. 2
- [3] Anish Athalye, Logan Engstrom, Andrew Ilyas, and Kevin Kwok. Synthesizing robust adversarial examples. In *Inter*national Conference on Machine Learning (ICML), 2018. 4
- [4] Melika Behjati, Seyed-Mohsen Moosavi-Dezfooli, Mahdieh Soleymani Baghshah, and Pascal Frossard. Universal adversarial attacks on text classifiers. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7345–7349, 2019. 3, 5
- [5] Nicholas Carlini and Hany Farid. Evading deepfake-image detectors with white- and black-box attacks. *arXiv preprint arXiv:2004.00622*, 2020. 1, 3
- [6] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *Symposium on Security* and Privacy, pages 39–57, 2017. 3, 4
- [7] Shuyu Cheng, Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Improving black-box adversarial attacks with a transfer-based prior. In *Advances in Neural Information Processing Systems (NIPS)*, pages 10934–10944, 2019. 3
- [8] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Conference on Computer Vision*

and Pattern Recognition (CVPR), pages 1251–1258, 2017. 2, 6

- [9] Azat Davletshin. https://github.com/ntech-lab/deepfakedetection-challenge. 2, 6
- [10] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. RetinaFace: Single-shot multi-level face localisation in the wild. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. 6
- [11] Brian Dolhansky, Joanna Bitton, Ben Pflaum, Jikuo Lu, Russ Howes, Menglin Wang, and Cristian Canton Ferrer. The DeepFake Detection Challenge (DFDC) dataset. arXiv preprint arXiv:2006.07397, 2020. 1, 2
- [12] Brian Dolhansky, Russ Howes, Ben Pflaum, Nicole Baram, and Cristian Canton Ferrer. The DeepFake Detection Challenge (DFDC) preview dataset. arXiv preprint arXiv:1910.08854, 2019. 2, 6
- [13] Apurva Gandhi and Shomik Jain. Adversarial perturbations fool deepfake detectors. arXiv preprint arXiv:2003.10596, 2020. 1, 3
- [14] Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In International Conference on Learning Representations (ICLR), 2015. 1, 3
- [15] Cui Hao. https://github.com/cuihaoleo/kaggle-dfdc. 2, 6
- [16] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Conference on Computer Vision and Pattern Recognition* (*CVPR*), pages 4401–4410, 2019. 2
- [17] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. arXiv preprint arXiv:1607.02533, 2016. 4
- [18] Jian Li, Yabiao Wang, Changan Wang, Ying Tai, Jianjun Qian, Jian Yang, Chengjie Wang, Ji-Lin Li, and Feiyue Huang. DSFD: dual shot face detector. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. 6
- [19] Yuezun Li and Siwei Lyu. Exposing deepfake videos by detecting face warping artifacts. In *Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 46–52, 2019. 2
- [20] Hong Liu, Rongrong Ji, Jie Li, Baochang Zhang, Yue Gao, Yongjian Wu, and Feiyue Huang. Universal adversarial perturbation via prior driven uncertainty approximation. In *International Conference on Computer Vision (ICCV)*, 2019. 3
- [21] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into transferable adversarial examples and blackbox attacks. In *International Conference on Learning Representations (ICLR)*, 2017. 3
- [22] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. Universal adversarial perturbations. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3, 5
- [23] Konda Reddy Mopuri, Utsav Garg, and R Venkatesh Babu. Fast feature fool: A data independent approach to universal adversarial perturbations. *arXiv preprint arXiv:1707.05572*, 2017. 3

- [24] Paarth Neekhara, Shehzeen Hussain, Malhar Jere, Farinaz Koushanfar, and Julian McAuley. Adversarial deepfakes: Evaluating vulnerability of deepfake detectors to adversarial examples. *arXiv preprint arXiv:2002.12749*, 2020. 1, 2, 3, 4, 6
- [25] Paarth Neekhara, Shehzeen Hussain, Prakhar Pandey, Shlomo Dubnov, Julian McAuley, and Farinaz Koushanfar. Universal adversarial perturbations for speech recognition systems. In *Interspeech*, 2019. 3, 5
- [26] Yuval Nirkin, Yosi Keller, and Tal Hassner. FSGAN: Subject agnostic face swapping and reenactment. In *Conference* on *Computer Vision and Pattern Recognition (CVPR)*, pages 7184–7193, 2019. 2
- [27] Nicolas Papernot, Patrick McDaniel, and Ian Goodfellow. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. arXiv preprint arXiv:1605.07277, 2016. 3
- [28] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In ACM Asia Conference on Computer and Communications Security, pages 506–519, 2017. 3
- [29] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In ACM on Asia Conference on Computer and Communications Security, 2017. 3
- [30] Nicolas Papernot, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Z Berkay Celik, and Ananthram Swami. The limitations of deep learning in adversarial settings. In *European Symposium on Security and Privacy (EuroS&P)*, 2016. 3
- [31] Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. Distillation as a defense to adversarial perturbations against deep neural networks. In Symposium on Security and Privacy (SP), pages 582–597, 2016.
- [32] Nicolas Rahmouni, Vincent Nozick, Junichi Yamagishi, and Isao Echizen. Distinguishing computer graphics from natural images using convolution neural networks. In Workshop on Information Forensics and Security (WIFS), 2017. 2
- [33] Konda Reddy Mopuri, Phani Krishna Uppala, and R Venkatesh Babu. Ask, acquire, and attack: Data-free uap generation using class impressions. In *European Conference* on Computer Vision (ECCV), pages 19–34, 2018. 3
- [34] Andreas Rossler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Niessner. FaceForensics++: Learning to detect manipulated facial images. In *International Conference on Computer Vision (ICCV)*, 2019.
 2
- [35] Selim Seferbekov. https://github.com/selimsef/dfdc_deepfake_ _challenge. 2, 3, 6
- [36] Yucheng Shi, Siyu Wang, and Yahong Han. Curls & whey: Boosting black-box adversarial attacks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [37] Dawn Song, Kevin Eykholt, Ivan Evtimov, Earlence Fernandes, Bo Li, Amir Rahmati, Florian Tramèr, Atul Prakash, and

Tadayoshi Kohno. Physical adversarial examples for object detectors. In USENIX Workshop on Offensive Technologies, 2018. 3

- [38] J.T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller. Striving for simplicity: The all convolutional net. In *International Conference on Learning Representations* (*ICLR*), 2015. 2
- [39] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013. 1, 3
- [40] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations (ICLR)*, 2014. 1, 3
- [41] Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. arXiv preprint arXiv:1905.11946, 2019. 6
- [42] Luisa Verdoliva. Media forensics and deepfakes: an overview. *arXiv preprint arXiv:2001.06564*, 2020. 1
- [43] Mika Westerlund. The emergence of deepfake technology: A review. *Technology Innovation Management Review*, 9:40– 53, 2019. 1
- [44] Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. Few-shot adversarial learning of realistic neural talking head models. In *International Conference* on Computer Vision (ICCV), pages 9459–9468, 2019. 2
- [45] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10):1499–1503, 2016. 6
- [46] Wen Zhou, Xin Hou, Yongjun Chen, Mengyun Tang, Xiangqi Huang, Xiang Gan, and Yong Yang. Transferable adversarial perturbations. In *European Conference on Computer Vision (ECCV)*, pages 452–467, 2018. 3