

Investigating the Impact of Pre-processing and Prediction Aggregation on the DeepFake Detection Task

Polychronis Charitidis Giorgos Kordopatis-Zilos Symeon Papadopoulos Ioannis Kompatsiaris

Information Technologies Institute, CERTH, Thessaloniki, Greece

{charitidis, georgekordopatis, papadop, ikom}@iti.gr

Abstract

Recent advances in content generation technologies (widely known as DeepFakes) along with the online proliferation of manipulated media content render the detection of such manipulations a task of increasing importance. Even though there are many DeepFake detection methods, only a few focus on the impact of dataset preprocessing and the aggregation of frame-level to video-level prediction on model performance. In this paper, we propose a pre-processing step to improve the training data quality and examine its effect on the performance of DeepFake detection. We also propose and evaluate the effect of video-level prediction aggregation approaches. Experimental results show that the proposed pre-processing approach leads to considerable improvements in the performance of detection models, and the proposed prediction aggregation scheme further boosts the detection efficiency in cases where there are multiple faces in a video.

1 Introduction

The latest advances in synthetic media manipulation have started posing new risks for society and democracy. Although the ability to generate or manipulate facial cues using artificial intelligence could have beneficial applications (Chandler, 2020; Hao, 2019) (e.g. art, video games, face anonymization, cinematography), there are several applications that are potentially harmful for individuals, communities and the society as a whole.

The term “DeepFake” initially referred to a deep learning-based method able to tamper media by swapping the face between two people. It appeared in 2017 when a machine learning algorithm was employed to transpose celebrity faces into porn videos. Apart from pornography, some of the most harmful usages of this technology include its use for online disinformation and financial fraud (Deeptrace, 2019).

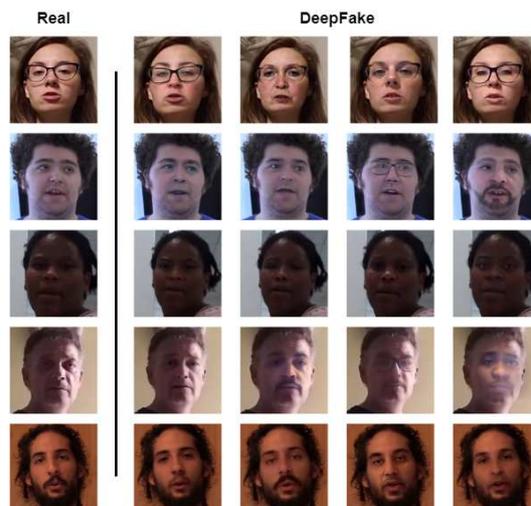


Figure 1: DFDC dataset examples (Kaggle, 2019)

However, “DeepFakes” have recently become synonymous with most types of facial and/or audio manipulation. Such manipulations typically include face swap, face generation from scratch, facial attribute manipulation, and facial expression manipulation/reenactment (Tolosana et al., 2020). Figure 1 illustrates a few example real faces from the DFDC dataset (Kaggle, 2019) and their corresponding DeepFake manipulations.

The harmful effects of DeepFakes are widely acknowledged in the research community, and significant effort has been invested lately for the detection of DeepFakes (Afchar et al., 2018; Rossler et al., 2019; Dang et al., 2020; Nguyen et al., 2019b; Li and Lyu, 2019; Cozzolino et al., 2018; Nguyen et al., 2019a; Güera and Delp, 2018; Sabir et al., 2019; Amerini et al., 2019). Several methods leverage recent advances in deep learning, like the ability to automatically learn rich features with Convolutional Neural Networks (CNNs). The detection problem is typically tackled by training a neural network in a supervised fashion to predict whether an input face is manipulated or not.

Deep learning approaches for DeepFake detection require the availability of large scale datasets. There are numerous such datasets, and the field is progressing rapidly. Besides academic contributions, even large companies like Google, contribute to DeepFake detection research by providing face manipulation datasets (Google AI Blog, 2019). Recently, AWS, Facebook, Microsoft, and the Media Integrity Steering Committee of the Partnership on AI launched the DeepFake Detection Challenge (DFDC) on Kaggle (Kaggle, 2019), offering a total prize of \$1,000,000, reflecting the importance that major stakeholders attribute to this problem.

Despite the rapid progress in DeepFake detection and dataset availability (Verdoliva, 2020), there has been very little focus on the pre-processing of training data and how this affects the performance of the resulting detection models. Pre-processing includes all transformations performed on the raw data before they are provided to a model for training or inference. In terms of videos, such transformations often include frame extraction, face detection, image normalization and resizing, and image augmentation.

In this work, we focus on face detection, an essential step for building accurate DeepFake detectors, as according to (Rossler et al., 2019), building detection models around a face results in higher accuracy compared to building models for whole images. A face detector with a large number of false positives will potentially generate a noisy dataset, and this might hurt the overall performance of the detection system. Consequently, a DeepFake detector’s performance depends heavily on the performance of the face detection model. After experimenting with various face detectors, we noticed many cases of false positive detections. To alleviate this issue, we propose a simple yet efficient pre-processing step that is applied after face extraction and can effectively remove a large number of falsely detected faces. Furthermore, we utilize the information from the proposed approach and develop a prediction aggregation scheme to improve the final video-level prediction and investigate its impact on the detection model compared to other aggregation baselines. Both the proposed data pre-processing scheme and the prediction aggregation improve detection performance significantly on three benchmark datasets.

The work presented in this paper is conducted

in the context of the *WeVerify* project (Marinova et al., 2020), which aims to build an open-source platform that engages communities and citizen journalists alongside newsroom and freelance journalists for collaborative and decentralised content verification, tracking, and debunking.

2 Related Work

2.1 DeepFake detection

Since the advent of deep learning, most classification tasks employ deep learning architectures that usually outperform traditional machine learning models. Following this trend, recent DeepFake detection approaches are based on deep learning networks for the detection of manipulated media.

Afchar et al. (2018) present two simple architectures with few layers that exploit mesoscopic features. Meso-4 has four layers of convolutions and pooling followed by a fully-connected layer for classification. Instead, MesoInception-4 is based on a simple variant of the inception module (Szegedy et al., 2015). XceptionNet (Chollet, 2017) is proposed as an efficient DeepFake detection architecture (Rossler et al., 2019). The same work shows that very deep general-purpose networks outperform shallow CNNs, like MesoInception-4, in the detection task. Dang et al. (2020) include an attention mechanism to their proposed architecture, which outperformed XceptionNet, when trained on a deepfake dataset they created with various deepfake manipulations. Nguyen et al. (2019b) presented a capsule-network that requires fewer parameters to train compared with the very deep networks and outperforms shallow nets like MesoInception-4. Li and Lyu (2019) present an approach that exploits artifacts from generated faces with limited resolution, potentially limiting the applicability of the method to specific generator models. Other works (Yang et al., 2019; Li et al., 2018) detect manipulations utilizing head pose and eye blinking information, respectively; these methods are only evaluated on specific types of manipulations. A more general approach is presented by Cozzolino et al. (2018), where an autoencoder-based architecture is proposed to adapt to new manipulations using just a few examples. This method outperforms XceptionNet in many types of manipulations. The same approach is used by Nguyen et al. (2019a) combining the detection and segmentation tasks to further assist the learning process. This

method is promising but as the previous one, requires manipulation masks for training. In (Güera and Delp, 2018), a convolutional Long Short Term Memory (LSTM) network is used to exploit temporal dependencies and provide video-level predictions. This work also investigates the impact of the numbers of frames on the final prediction and it reports no further performance gains using more than 40 frames. A recurrent convolutional model was also proposed by Sabir et al. (2019) with competitive results but limited evaluations. Amerini et al. (2019) leverage the optical flow to exploit temporal discrepancies among frames and show that optical flow is predictive of DeepFake manipulations.

2.2 DeepFake datasets

Of the four major categories of DeepFake manipulations, we mainly present the datasets related to face swapping, facial expression manipulation and facial attribute manipulation, as they are the ones we focus on in this work. UADFV (Yang et al., 2019) is an initial small-scale dataset employing face swapping. The authors in (Korshunov and Marcel, 2018) present the DeepFake-TIMIT dataset. This consists of 620 fake videos created using a GAN-based face swapping algorithm. FaceForensics++ (Rossler et al., 2019) is a popular DeepFake dataset that contains 1000 real videos from YouTube. This dataset provides fake videos using face swapping and face reenactment manipulation techniques. The Google/Jigsaw also contributed to the FaceForensics++ dataset with the DeepFake detection dataset (DFD) (Google AI Blog, 2019). The Celeb-DF (Li et al., 2020) dataset aims to provide face swapping videos of better visual qualities, as previous databases exhibit low visual quality with many visible artifacts. Celeb-DF consists of 408 real videos from YouTube and 795 fake videos. More recently, the DeepFake Detection Challenge (DFDC) (Kaggle, 2019; Dolhansky et al., 2019) first released a preview dataset consisting of 1131 real videos from 66 paid actors, and 4113 fake videos. The complete DFDC dataset was released on the 11th of December 2019, containing approximately 20,000 real videos and 100,000 fakes. The authors in (Jiang et al., 2020) present another large-scale DeepFake dataset, including 10,000 fake videos, built using 100 actors.

2.3 Dataset pre-processing

Although many DeepFake detection works employ data pre-processing, there is very little discussion on the impact of this step on the final detection model. Rossler et al. (2019) report that extracting the face region of a video frame, instead of using the whole frame as input to the deep learning model, yields better detection. Numerous works adopt this approach and use face detection libraries to pre-process the video frames and extract face images. There are many face detection works available (King, 2009; Zhang et al., 2016; Bazarevsky et al., 2019; Deng et al., 2019) and multiple implementations of them, which mainly differ in accuracy, detection speed, and setting availability (e.g., batched detection, detection threshold). In addition to face detection, other works adopt face tracking or face alignment approaches. For example, the authors in (Sabir et al., 2019) examine the impact of explicit alignment using facial landmarks and implicit alignment that uses a Spatial Transformer Network (STN) (Jaderberg et al., 2015) and find that landmark alignment improves the DeepFake detection performance of their presented architecture.

3 Baseline DeepFake Detection Pipeline

Figure 2 illustrates a baseline approach for building a video DeepFake detection model. For the training and evaluation of a detection model, one or more of the DeepFake datasets (Section 2.2) can be selected. To transform the raw videos into a format that can be used by deep learning architectures, we apply the three steps depicted in the pre-processing block in Figure 2.

In the first pre-processing step, we extract video frames by applying uniform sampling based on the video duration. Next, we extract the face regions detected in each video frame, as recommended by (Rossler et al., 2019). In this step, the face detectors listed in 2.3 can be used to detect faces and return the face coordinates (in the form of bounding boxes) in the corresponding frames. Note that in this step, a common practice is to include some background regions along with the face, as in Rossler et al. (2019). Thus, we multiply the face bounding boxes by a factor of 1.3. The main reason for this is to enable deep learning models to detect resolution inconsistencies or other discrepancies between the face and its surroundings.

The final pre-processing step includes the fol-

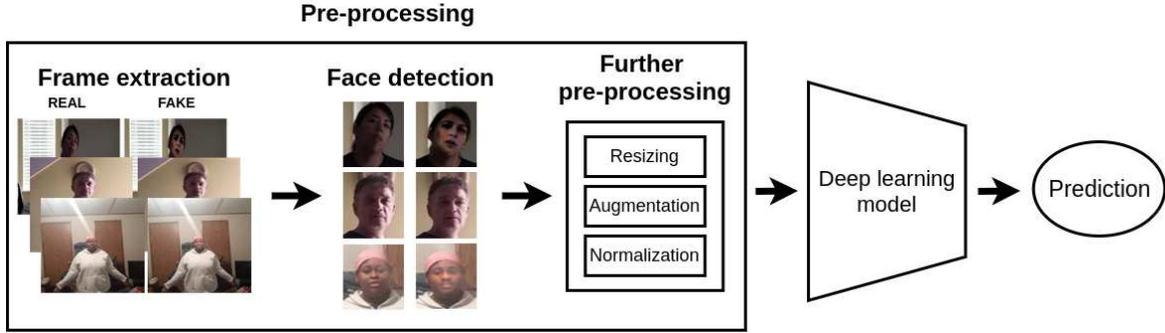


Figure 2: Baseline DeepFake detection pipeline

lowing transformations: (a) input face image resizing, (b) image augmentation, (c) image normalization, which generally refers to the translation of the RGB pixel values from $[0-255]$ to $[0-1]$. Such transformations can prevent overfitting in training and generally lead to more robust classifiers.

Using this baseline detection setup, we train state-of-the-art DeepFake detection architectures. The employed architectures operate at image-level, meaning that we train the models with individual images, and we optimize the model to detect manipulations on them. An additional post-processing step is required to aggregate the per image predictions to a single prediction for the entire video during inference. We experiment with different post-processing aggregation methods for video-level predictions, and we investigate an aggregation approach that utilizes the concept of the connected components according to the proposed pre-processing step (Section 4).

The issue with the pre-processing pipeline of Figure 2 is that it depends on the accurate face detection for the generation of face images. Having experimented with multiple face detectors, we noticed that the amount of false positives is higher than expected. Of course, one could fine-tune the detection settings (e.g., increase the detection confidence threshold) to minimize false detections, but this process is time-consuming, depends on the examined dataset and runs the risk of removing correctly detected faces. Using the default settings of a publicly available implementation for face detection Zhang et al. (2016), we extracted the face regions from some random DFDC videos. The results are presented in Figure 3. One may notice that among correctly detected faces, there are cases where the face detector failed to detect a human face. False detections usually include random shapes, various human body parts (e.g.,



Figure 3: Extracted face detection regions from random DFDC videos (rows). Among the detected faces, there are several cases of false positive detection.

hands, neck) and regions with a small proportion of the face displayed. Additionally, we empirically found that false detections are usually not consistent across the duration of videos, meaning that they do not appear in every extracted video frame.

However, false detections can lead to noisy training data, which could impact the DeepFake detection performance. To address this issue, we propose an additional pre-processing step to clean the dataset by removing detected images that do not contain faces.

4 Proposed Pre-processing Approach

Here, we describe the proposed pre-processing step for removing false face detections, which we apply after the face detection step (Figure 2).

4.1 Method description

The main objective of the proposed step is to generate clusters with correct and incorrect detections, i.e., face and non-face images, so as to remove the latter. We have noticed that false detections occur randomly throughout the video, and they are not repeated in every frame. Furthermore, faces should be present throughout the video, which is valid for several DeepFake datasets. Therefore, clusters formed by incorrect detections should

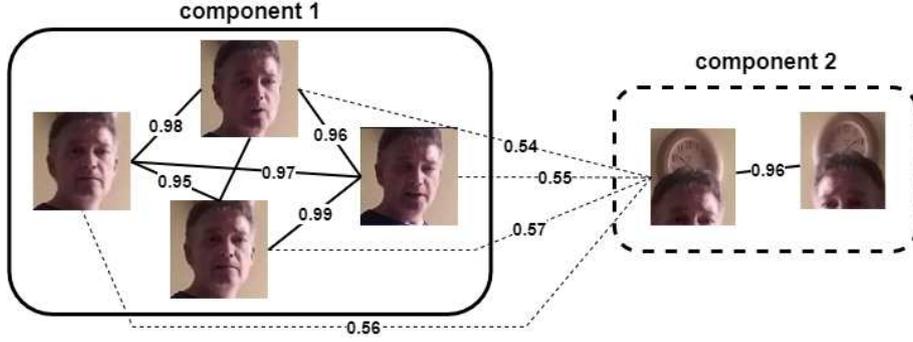


Figure 4: The proposed pre-processing step. Images that are similar to each other (solid lines) form connected components. Dashed lines connect images with low similarity. For simplicity, we show only a subset of possible connections among faces. Components of small size (cf. section 4.1) (dashed border) are removed.

have a much smaller size than clusters of correctly detected faces.

To this end, we employ a face recognition model, based on the work by Schroff et al. (2015), in order to compute facial embeddings for the extracted images. Embeddings encode the facial information in D -dimensional vectors. The employed architecture extracts embeddings with $D = 512$. To this end, we calculate the similarity between the detected faces based on the dot product of the corresponding embedding vectors. More formally, let $i, j \in \{1, 2, \dots, K\}$ where K is the number of detected faces in N extracted frames, then the similarity between the i -th and j -th detected face is defined as:

$$s(i, j) = f_w(i)^\top \cdot f_w(j) \quad (1)$$

where $f_w(\cdot)$ is the embedding function that maps an arbitrary face image to the facial embedding space \mathbb{R}^D , and $s(\cdot, \cdot)$ is the function that assesses the similarity between two face images.

Then, we utilize the similarity information between the detected faces to generate connected components on a *face graph*. The graph nodes correspond to the detected face images. For the face graph formation, two nodes i, j are connected with an edge if their similarity is greater than a predefined threshold θ , i.e., $s(i, j) > \theta$. We experimentally validate the impact of the similarity threshold θ in evaluation datasets in Section 5.5.

Figure 4 illustrates this process. Nodes with similarity greater than $\theta = 0.8$ are connected with each other (solid line); otherwise, there is no edge between nodes (dashed line). After this process is completed, if there are no false detections, we expect the number of the generated components to be equal to the number of distinct faces in a

video. In cases where there are false detections, these should form an independent component, i.e., Component 2 in Figure 4. As mentioned above, these components will usually contain fewer nodes in comparison to the ones formed by correct detections. This is valid in most cases because the face detectors produce false results in only a subset of the frames. We have empirically found that removing components with size less or equal than $N_F/2$, where N_F is the number of the sampled frames of a video containing at least one detected face, leads to better dataset quality. We qualitatively checked 50 random videos from the DFDC dataset. Figure 5 depicts some of the qualitative results. Furthermore we investigate the impact of the component size threshold in evaluation datasets in Section 5.5. In the provided example in Figure 4, let $N_F = 4$ ¹, then Component 2 with only 2 nodes is removed, and only images from component 1 are forwarded to the next step of the pre-processing pipeline.

4.2 Advantages and limitations

The main advantages of the proposed approach are that it is a simple way to increase detection performance, and it is very fast. Also, it can be used on top of any existing face detection process, in combination with any available face detection library.

Furthermore, the facial embedding information can be utilized in order to separate the detected faces to clusters of different persons. This functionality is not available by most face detection implementations. This is particularly useful for making separate predictions per face (person), es-

¹Note that in this particular example, some of the detected faces may come from the same frame and this is why N_F is smaller than the number of nodes in the face graph.

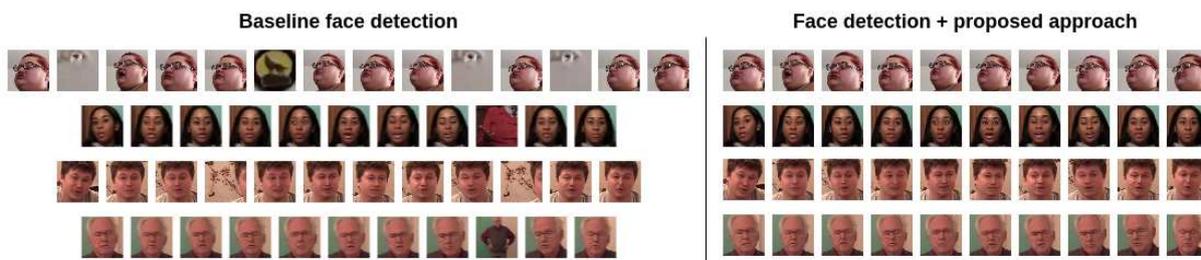


Figure 5: Qualitative results from DFDC videos. Faces extracted with baseline preprocessing (left). Faces extracted adding the proposed preprocessing step (right). Frames sampled with 1 frame per second.

pecially for cases where there is only one manipulated face among many in a video. This technique is invariant to face movements and it can accurately track multiple moving faces and form separate face clusters.

Our approach has two main limitations. The first is that it assumes that faces are present throughout the duration of a video. So, in cases where a person appears only for a small fraction of the video duration, this face will likely be considered a false detection and consequently will be removed. Even though this is generally not the case in the DeepFake datasets, it can be encountered in several real-world online manipulated videos. The second limitation is that we consider clusters of small sizes to be false detections, and although this is usually the case, there are cases where the face detector can make the same incorrect detections in every extracted frame. In that case, the resulting cluster will be large and hence our approach will incorrectly assume that it corresponds to a correctly detected face.

5 Experimental Study

5.1 Pre-processing setup

To examine the impact of pre-processing on the detection accuracy, we evaluate two different pre-processing approaches. The first approach, which we denote as baseline, is based on the pipeline described in Section 3. The second approach additionally contains the proposed pre-processing step, as described in Section 4, after the face detection.

For both pre-processing approaches, we use the method proposed in Zhang et al. (2016) to detect faces, setting the face detection threshold to 0.9, as a good trade-off between reducing false detections and not increasing false negatives. We also expand the size of a detected bounding box by a factor of 1.3, as reported in (Rossler et al., 2019). Additionally, for model training, we apply several

augmentations on the extracted images: horizontal and vertical flipping, random cropping, rotation, compression, Gaussian and motion blurring, and brightness, saturation, and contrast transformation². Also, we normalize the input RGB values to the [0-1] interval. For ImageNet pre-trained models, i.e., XceptionNet and EfficientNet-B4, we employ the normalization scheme they have been trained with, including the channel-wise subtraction of the ImageNet dataset mean, followed by division with the ImageNet dataset standard deviation. In the rare case where no faces are detected in a video during training, we do not further consider the video, while during the inference phase, we manually set the prediction to 0.5.

5.2 Training setup

We use the DFDC dataset for model training. The dataset contains approximately 20,000 real videos and 100,000 fakes. Of those, we use 1000 videos for validation, 1202 for testing, and the rest for training. The validation and test sets have equal number of real and fake videos. We also generate two face datasets with these videos using both the baseline and the proposed pre-processing approach. For the proposed training dataset generation, we apply our pre-processing approach to all frames of the video in order to extract faces. To deal with the class imbalance, we sample 16 of these faces from the real videos and only 4 from the fake ones during batch generation. We randomly sample the same amount of frames every epoch from each video to increase training diversity. For the baseline pre-processing dataset, we sample the same number of detected faces.

We experiment with three different deep learning architectures: MesoInception-4 (Afchar et al., 2018), which was designed specifically for Deep-

²For implementation, we made use of the `imgaug` Python library, <https://imgaug.readthedocs.io/en/latest/index.html>

Fake detection, XceptionNet (Chollet, 2017), which outperforms other deep networks (Rossler et al., 2019), and EfficientNet (Tan and Le, 2019) (the EfficientNet-B4 variant), which achieves state-of-the-art performance in most image classification tasks. To adapt the last two architectures to the DeepFake setting, we employ the corresponding backbone networks and add two fully connected layers with 512 and 1 neurons, respectively. We use the sigmoid activation function in the final layer for classification. Except for MesoInception-4, the other models are initialized using the ImageNet pre-trained weights (Russakovsky et al., 2015). For training, we use the Adam optimizer (Kingma and Ba, 2014) and minimize the Log loss error. Note that the training process operates at the image and not at video level. We train the networks for 10 epochs with a learning rate of 10^{-4} , and with batch size of 84 for MesoInception-4 and 16 for XceptionNet and EfficientNet. We select the model with the best validation error for each architecture.

5.3 Evaluation setup

To examine the impact of the different pre-processing approaches on the trained detection models we evaluate them on the Celeb-DF (Li et al., 2020), FaceForensics++ (Rossler et al., 2019) datasets and DFDC test datasets. Celeb-DF consists of 408 real and 795 fake videos. FaceForensics++ consists of 1000 real videos and 4000 fake videos. To balance datasets, we randomly subsample the majority class in the case of Celeb-DF. For FaceForensics++, we use 1000 real videos and only 1000 fake videos with DeepFakes manipulation, ignoring the other three manipulations (Face2Face, FaceSwap, Neural Textures). Since the training dataset and DFDC test dataset originate from the same distribution of videos, i.e., the same face manipulations were used, we expect detection performance in this dataset to be significantly better in comparison to the other two evaluation datasets.

For evaluation, we extract 4 frames per second and we run experiments using the detection models that were trained with the baseline and proposed pre-processing approaches. For the proposed pre-processing approach, we make separate predictions for every detected face image in a video. To aggregate these predictions, we consider four approaches. a) averaging the individ-

ual predictions (*Avg*), b) taking the median prediction (*Median*), c) taking the maximum prediction (*Max*), d) averaging predictions per face using the component information from the pre-processing step and taking the maximum prediction among faces (*Face*). To measure detection performance, we report the aggregated video-level Log loss error and accuracy for each setting.

5.4 Experimental results

Table 1(a) illustrates the Log loss error of the three benchmarked models on every evaluation dataset. Note that the error for a dummy classifier that always predicts 0.5 for each video is 0.693. For the case of the Celeb-DF dataset, we notice that EfficientNet-B4 achieves the best performance among all models. MesoInception-4 marginally surpasses the performance of the dummy classifier. This indicates that shallow architectures are not suitable for the DeepFake detection task. Models trained with the proposed pre-processing approach outperform the baseline by a significant margin. For example, the performance gain of the models with *Avg* aggregation ranges between 0.033 and 0.055 in terms of Log loss error, corresponding to 5-11% relative performance increase. In terms of the aggregation methods, the *Face* and *Avg* methods achieve the best results exhibiting similar performance, with the former achieving marginally better results. This is expected because videos in the Celeb-DF dataset contain one face throughout the video, and as a result, both approaches average the predictions, except for some cases where the pre-processing approach detects multiple components. The *Median* approach is slightly worse than the previous two, and *Max* achieves the worst performance.

Similar conclusions apply for the FaceForensics++ dataset. EfficientNet-B4 with the *Avg* aggregation achieves the best performance. Once again, it is clear that pre-processing is beneficial for the performance of the detection models, even more than in the Celeb-DF dataset. In almost all model-aggregation combinations, the Log loss error drops considerably when the network is trained with the proposed approach compared to the baseline. For example, the performance gain of the models with *Avg* ranges between 0.035 and 0.067, corresponding to a 5-12% relative performance increase. Also, *Face* and *Avg* aggregations demonstrate similar performance, which can be attributed

Table 1: Log loss error, Accuracy (%) and F1 score of three architectures on three DeepFake datasets. For comparison, note that the Log loss error of the dummy classifier (always predict 0.5) is 0.693.

Model	Pre-processing	Celeb-DF				FaceForensics++				DFDC test			
		Avg	Med.	Max	Face	Avg	Med.	Max	Face	Avg	Med.	Max	Face
MesoInc4	baseline	0.677	0.689	0.782	0.679	0.670	0.680	0.769	0.672	0.484	0.491	0.540	0.463
	proposed	0.644	0.657	0.791	0.642	0.635	0.639	0.705	0.636	0.420	0.441	0.503	0.401
Xception	baseline	0.562	0.573	0.601	0.561	0.584	0.598	0.613	0.585	0.354	0.364	0.402	0.333
	proposed	0.520	0.530	0.549	0.522	0.543	0.546	0.580	0.544	0.312	0.324	0.382	0.292
EffNet-B4	baseline	0.510	0.519	0.569	0.508	0.568	0.578	0.601	0.566	0.213	0.223	0.310	0.198
	proposed	0.463	0.495	0.558	0.453	0.497	0.515	0.621	0.499	0.195	0.243	0.320	0.173

(a) Log loss error

Model	Pre-processing	Celeb-DF				FaceForensics++				DFDC test			
		Avg	Med.	Max	Face	Avg	Med.	Max	Face	Avg	Med.	Max	Face
MesoInc4	baseline	67.8	67.5	61.2	67.2	66.7	66.4	60.2	66.4	75.8	75.7	64.0	76.9
	proposed	69.3	69.1	61.2	69.6	67.2	67.0	59.8	66.8	76.6	76.5	64.4	78.4
Xception	baseline	76.8	76.7	63.0	76.8	75.4	75.1	67.8	75.2	85.5	85.0	69.1	85.8
	proposed	78.5	78.2	66.3	78.6	77.2	76.0	68.1	77.0	86.1	87.7	70.1	87.8
EffNet-B4	baseline	81.8	81.4	72.1	81.6	78.4	78.1	69.1	78.9	92.1	92.1	84.3	93.8
	proposed	83.0	82.9	71.0	83.1	81.2	80.8	70.2	81.0	94.2	94.5	84.9	96.3

(b) Accuracy

Model	Pre-processing	Celeb-DF				FaceForensics++				DFDC test			
		Avg	Med.	Max	Face	Avg	Med.	Max	Face	Avg	Med.	Max	Face
MesoInc4	baseline	64.2	63.0	59.5	64.6	63.9	64.2	58.2	64.2	74.2	74.0	63.3	76.2
	proposed	67.0	66.1	62.2	67.6	65.2	64.9	58.7	66.1	77.1	76.5	65.4	78.8
Xception	baseline	75.5	74.9	60.1	75.3	73.9	72.8	67.2	73.5	86.2	85.5	69.9	86.3
	proposed	77.2	76.9	62.3	77.6	75.8	74.2	66.9	75.2	87.2	86.8	73.1	88.2
EffNet-B4	baseline	82.7	81.4	75.1	82.5	77.6	76.1	73.2	77.0	93.0	92.5	83.5	94.3
	proposed	83.1	82.5	74.0	84.1	79.2	78.8	73.2	78.4	95.1	94.5	83.9	96.5

(c) F1 score

to the fact that the FaceForensics++ dataset contains one face per video as well.

The Log loss error of all models on the DFDC test is significantly smaller compared to other datasets. This is expected because the test set contains videos with the same manipulations as the training set. It is noteworthy that the models trained with the proposed pre-processing approach consistently outperform those trained with the baseline. In this case, models with the proposed pre-processing score 12-13% better. Additionally, the *Face* aggregation scheme outperforms the *Avg* by a clear margin. This can be attributed to the fact that the DFDC dataset contains several videos with more than one person. Also, only one of the faces in these videos are manipulated, which would render *Avg* an almost random way of aggregating individual predictions.

Table 1(b) displays the results in terms of accuracy. We derive similar conclusions as in the case of Log loss error. In general, accuracy is improved when training the models with the proposed pre-processing pipeline. The average gain for *Face* and *Avg* aggregation methods is approx-

imately 2%. The *Face* aggregation scores slightly better than *Avg* for the Celeb-DF and the opposite for the FaceForensics++ dataset. The *Face* aggregation method outperforms *Avg* by a larger margin for the case of DFDC test dataset.

For completeness, we also report the macro-average F1 score in Table 1(c). Macro-average F1 improved when training the models with the proposed pipeline. Reporting the best aggregation approach per dataset, we observe a 2.3% performance increase with the proposed pre-processing for the Celeb-DF dataset. This value is 1.6% and 2.2% for the FaceForensics++ and DFDC test datasets respectively.

5.5 Impact of selected thresholds

In this section, we investigate the impact of the similarity threshold and component size threshold on detection performance. We pre-process evaluation datasets with the proposed pre-processing approach using different similarity thresholds and employ the detection models to evaluate them. We also investigate the impact of the component size threshold on the model’s performance testing

Table 2: Log loss error of the EfficientNet-B4 architecture with different similarity thresholds and varying component sizes on three DeepFake datasets

Dataset	Similarity threshold			Size threshold		
	0.7	0.8	0.9	$N_F/4$	$N_F/2$	$3N_F/4$
Celeb-DF	0.480	0.453	0.445	0.453	0.453	0.475
FaceForensics++	0.525	0.499	0.506	0.510	0.499	0.544
DFDC test	0.204	0.173	0.194	0.182	0.173	0.230

three different threshold values on the evaluation datasets. We experiment only with the EfficientNet model with the *face* aggregation method.

Table 2 displays the Log loss error of our model for three similarity thresholds on three test datasets. It is evident that using 0.8 similarity threshold yields the best results for the case of FaceForensics++ and DFDC test datasets. The best performance in Celeb-DF is achieved with 0.9 similarity threshold. Generally, the performance difference between 0.8 and 0.9 similarity thresholds is marginal. On the other hand, the performance drop is more significant when the similarity threshold is lower ($\theta \leq 0.7$). This is expected as lowering the similarity threshold leads to generated components with more dissimilar faces. This means that more false detections are connected with true detections to form a component, limiting the dataset cleaning capabilities.

Table 2 also shows the experimental results when applying the proposed pre-processing with different minimum component size threshold. The similarity threshold for these experiments is 0.8. In the cases where all components are removed for a video, then we set the final video prediction to 0.5. When we increase the minimum size threshold to $3N_F/4$, then the evaluation error is higher. This happens because there are cases in the dataset where a component with correctly detected faces is removed, leading to no generated components for a particular video. On the other hand, we observe a measurable increase in the error when the size threshold decreases to $N_F/4$. For the case of the Celeb-DF dataset, the error remains the same as the initial one, meaning that there is no change in the detected components. In the other two evaluation datasets, a small increase in the error is observed, owing to the fact that the proposed process fails to remove some incorrectly detected components in certain videos.

6 Discussion and Conclusions

In this work, we studied the impact of dataset pre-processing and prediction aggregation on the DeepFake detection problem. We proposed a pre-processing step that improves the training dataset quality. We found that pre-processing has a considerable impact on the detection task and boosts model performance by improving the quality of the generated training set. Using the information derived from this approach, we experimented with an aggregation approach that outperforms baselines when multiple faces appear in a video.

A noteworthy observation of our work is the lack of detection generalization in unseen manipulations. It is apparent that when the detection models are trained with a dataset of specific facial manipulations, then the performance on other datasets with different manipulations is poor. This means that research efforts have to focus on the generalization of the detection models, especially now that new manipulation methods appear by the day. Additionally, the broad definition of DeepFake manipulations that includes all kinds of facial manipulation and face generation makes the problem even harder.

In the future, we plan to further optimize the performance of our Deepfake Detection pipeline and deploy it as a web service and make it accessible to journalists and citizens through the web media forensics tool (Zampoglou et al., 2016) and the InVID-WeVerify verification plugin (Teyssou et al., 2017). In that way, we will explore both challenging real-world examples of DeepFakes (new types of manipulation, different compression schemes, etc.) and challenges related to user experience and explainability of the results.

Acknowledgments: This work has been supported by the WeVerify (contract nr. 825297) and AI4Media (contract nr. 951911) projects, partially funded by the European Commission.

References

- Darius Afchar, Vincent Nozick, Junichi Yamagishi, and Isao Echizen. 2018. Mesonet: a compact facial video forgery detection network. In *2018 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–7. IEEE.
- Irene Amerini, Leonardo Galteri, Roberto Caldelli, and Alberto Del Bimbo. 2019. Deepfake video detection through optical flow based cnn. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0.
- Valentin Bazarevsky, Yury Kartynnik, Andrey Vakunov, Karthik Raveendran, and Matthias Grundmann. 2019. Blazeface: Sub-millisecond neural face detection on mobile gpus. *arXiv preprint arXiv:1907.05047*.
- Simon Chandler. 2020. [Why deepfakes are a net positive for humanity](#).
- François Chollet. 2017. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1251–1258.
- Davide Cozzolino, Justus Thies, Andreas Rössler, Christian Riess, Matthias Nießner, and Luisa Verdoliva. 2018. Forensictransfer: Weakly-supervised domain adaptation for forgery detection. *arXiv preprint arXiv:1812.02510*.
- Hao Dang, Feng Liu, Joel Stehouwer, Xiaoming Liu, and Anil K Jain. 2020. On the detection of digital face manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5781–5790.
- Deeptrace. 2019. [The State Of Deepfakes: Landscape, Threats and Impact](#).
- Jiankang Deng, Jia Guo, Yuxiang Zhou, Jinke Yu, Irene Kotsia, and Stefanos Zafeiriou. 2019. Retinaface: Single-stage dense face localisation in the wild. *arXiv preprint arXiv:1905.00641*.
- Brian Dolhansky, Russ Howes, Ben Pflaum, Nicole Baram, and Cristian Canton Ferrer. 2019. The deepfake detection challenge (dfdc) preview dataset. *arXiv preprint arXiv:1910.08854*.
- Google AI Blog. 2019. [Contributing Data to Deepfake Detection Research](#).
- David Güera and Edward J Delp. 2018. 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.
- Karen Hao. 2019. [Deepfakes could anonymize people in videos while keeping their personality](#).
- Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. 2015. Spatial transformer networks. In *Advances in Neural Information Processing Systems*, pages 2017–2025.
- Liming Jiang, Ren Li, Wayne Wu, Chen Qian, and Chen Change Loy. 2020. Deepforensics-1.0: A large-scale dataset for real-world face forgery detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Kaggle. 2019. [Deepfake detection challenge](#).
- Davis E. King. 2009. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10:1755–1758.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Pavel Korshunov and Sébastien Marcel. 2018. Deepfakes: a new threat to face recognition? assessment and detection. *arXiv preprint arXiv:1812.08685*.
- Yuezun Li, Ming-Ching Chang, and Siwei Lyu. 2018. In icu 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.
- Yuezun Li and Siwei Lyu. 2019. Exposing deepfake videos by detecting face warping artifacts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 46–52.
- Yuezun Li, Xin Yang, Pu Sun, Honggang Qi, and Siwei Lyu. 2020. 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.
- Zlatina Marinova, Jochen Spangenberg, Denis Teyssou, Symeon Papadopoulos, Nikos Sarris, Alexandre Alaphilippe, and Kalina Bontcheva. 2020. Weverify: Wider and enhanced verification for you project overview and tools. In *2020 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, pages 1–4. IEEE.
- Huy H Nguyen, Fuming Fang, Junichi Yamagishi, and Isao Echizen. 2019a. Multi-task learning for detecting and segmenting manipulated facial images and videos. *arXiv preprint arXiv:1906.06876*.
- Huy H Nguyen, Junichi Yamagishi, and Isao Echizen. 2019b. Capsule-forensics: Using capsule networks to detect forged images and videos. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2307–2311. IEEE.
- Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. 2019. Faceforensics++: Learning to detect manipulated facial images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1–11.

- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252.
- Ekraam Sabir, Jiabin Cheng, Ayush Jaiswal, Wael AbdAlmageed, Iacopo Masi, and Prem Natarajan. 2019. Recurrent convolutional strategies for face manipulation detection in videos. *Interfaces (GUI)*, 3:1.
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 815–823.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9.
- Mingxing Tan and Quoc Le. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114.
- Denis Teysou, Jean-Michel Leung, Evlampios Apostolidis, Konstantinos Apostolidis, Symeon Papadopoulos, Markos Zampoglou, Olga Papadopoulou, and Vasileios Mezaris. 2017. The InVID plug-in: web video verification on the browser. In *Proceedings of the first International Workshop on Multimedia Verification*, pages 23–30.
- Ruben Tolosana, Ruben Vera-Rodriguez, Julian Fierrez, Aythami Morales, and Javier Ortega-Garcia. 2020. Deepfakes and beyond: A survey of face manipulation and fake detection. *arXiv preprint arXiv:2001.00179*.
- Luisa Verdoliva. 2020. Media forensics and deepfakes: an overview. *arXiv preprint arXiv:2001.06564*.
- Xin Yang, Yuezun Li, and Siwei Lyu. 2019. Exposing deep fakes using inconsistent head poses. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8261–8265. IEEE.
- Markos Zampoglou, Symeon Papadopoulos, Yianis Kompatsiaris, Ruben Bouwmeester, and Jochen Spangenberg. 2016. Web and social media image forensics for news professionals. In *Tenth international AAAI Conference on Web and Social Media*.
- Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. 2016. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10):1499–1503.