

Received August 6, 2020, accepted September 1, 2020, date of publication September 9, 2020, date of current version September 22, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.3023037

Fighting Deepfake by Exposing the Convolutional Traces on Images

LUCA GUARNERA^{1,2}, (Student Member, IEEE), OLIVER GIUDICE¹,
AND SEBASTIANO BATTIATO^{1,2}, (Senior Member, IEEE)

¹Department of Mathematics and Computer Science, University of Catania, 95124 Catania, Italy

²iCTLab s.r.l. Spinoff of University of Catania, 95124 Catania, Italy

Corresponding author: Luca Guarnera (luca.guarnera@unict.it)

This work was supported by iCTLab s.r.l. - Spin-off of University of Catania.

ABSTRACT Advances in Artificial Intelligence and Image Processing are changing the way people interacts with digital images and video. Widespread mobile apps like *FACEAPP* make use of the most advanced Generative Adversarial Networks (GAN) to produce extreme transformations on human face photos such gender swap, aging, etc. The results are utterly realistic and extremely easy to be exploited even for non-experienced users. This kind of media object took the name of Deepfake and raised a new challenge in the multimedia forensics field: the Deepfake detection challenge. Indeed, discriminating a Deepfake from a real image could be a difficult task even for human eyes but recent works are trying to apply the same technology used for generating images for discriminating them with preliminary good results but with many limitations: employed Convolutional Neural Networks are not so robust, demonstrate to be specific to the context and tend to extract semantics from images. In this paper, a new approach aimed to extract a Deepfake fingerprint from images is proposed. The method is based on the Expectation-Maximization algorithm trained to detect and extract a fingerprint that represents the Convolutional Traces (CT) left by GANs during image generation. The CT demonstrates to have high discriminative power achieving better results than state-of-the-art in the Deepfake detection task also proving to be robust to different attacks. Achieving an overall classification accuracy of over 98%, considering Deepfakes from 10 different GAN architectures not only involved in images of faces, the CT demonstrates to be reliable and without any dependence on image semantic. Finally, tests carried out on Deepfakes generated by *FACEAPP* achieving 93% of accuracy in the fake detection task, demonstrated the effectiveness of the proposed technique on a real-case scenario.

INDEX TERMS Deepfake detection, generative adversarial networks, multimedia forensics, image forensics.

I. INTRODUCTION

A digital image can be manipulated with many tools and software. Everyone with a glimpse of experience in using Photoshop or GIMP can forge photographs in order to change their contents, the semantics and - potentially - everything. However, this kind of forgery has been widely investigated throughout recent years and commercial tools with the ability to detect and describe them are also available [1], [2]. The possibility to detect forgeries made with Photoshop or similar tools are related to the experience of the image manipulator being able to hide any kind of unrealistic artifact.

The associate editor coordinating the review of this manuscript and approving it for publication was Aniello Castiglione¹.

Advances in Artificial Intelligence, and specifically, the advent of Generative Adversarial Networks (GAN) [3], enabled the creation and widespread of extremely refined techniques able to *attack* digital data, alter it or create its contents from scratch. These tools are able to obtain surprisingly realistic results leading to the birth of the Deepfake images phenomenon, or simply Deepfakes.

In general, a Deepfake is defined as a multimedia content synthetically modified or created through automatic (or barely controlled) machine learning models. Most state-of-the-art techniques are able to do the *face swap* from a source image/video to a target image/video. Recently, faces of showgirls, politicians, actors, TV presenters and many others have been the main protagonists of Deepfake attacks: one of the first example is the famous face swap of Jim Carrey on top



FIGURE 1. Examples of Deepfakes: (a) Jim Carrey's face transferred to Alison Brie's body, (b) Mr. Bean is Charlize Theron in a Deepfake version of J'adore commercial, (c) Jim Carrey instead of Jack Nicholson in *Shining* and (d) Tom Cruise Replaces Robert Downey Jr. in *Iron Man*.

of the the body of Alison Brie¹ (Figure 1a), or Mr. Bean and Charlize Theron in the Deepfake version of the commercial of J'adore² (Figure 1b), and again Jim Carrey instead of Jack Nicholson in *Shining*³ (Figure 1c), or Tom Cruise replacing Robert Downey Jr. in *Iron Man*⁴ (Figure 1d).

Deepfakes are not only involved in face-related tasks but they could be engaged to swap or generate realistic places, animals, object, etc. Indeed, this could bring disruptive innovation in many working areas, such as in the automotive industry or in architecture, since it is possible to generate a car or an apartment through dedicated GANs or in the film industry where it is possible, when necessary, to replace the face of a stuntman with an actor; but, on the other hand it could lead to serious social repercussions, privacy issues and major security concerns. For example, there are many Deepfake videos connected to the world of porn used to discredit famous actresses like Emma Watson or Angelina Jolie, or they can be used to spread disinformation and fake news. Moreover, the creation of Deepfakes is becoming extremely easy: widespread mobile apps like *FACEAPP*,⁵ are able to produce transformations on human faces such gender swap, aging, etc. The results are utterly realistic and extremely easy to produce even for non-experienced users with a few taps on their mobile phone.

It is clear that the Deepfake phenomenon raises a serious safety issue and it is absolutely necessary to create new techniques able to detect and counteract it [1], [4].

¹https://www.youtube.com/watch?v=SEar_6UtX9U

²<https://www.youtube.com/watch?v=gZVdPJhBkqg>

³<https://www.youtube.com/watch?v=JbzVhzNaTdl>

⁴<https://www.youtube.com/watch?v=iDM69UEyM3w>

⁵<https://www.faceapp.com/>

While detecting a Deepfake is difficult for humans, recent works have shown that they could be detected surprisingly easily by employing Convolutional Neural Networks (CNN) specifically trained on the task. However, CNN solutions presented till today, lack of robustness, generalization capability and explainability. They are extremely specific to the context in which they were trained and, being very deep, tend to extract the underlying semantics from images without inferring any unique fingerprint. A detailed discussion about such limits will be dealt with in the final part of the paper.

In order to find a unique fingerprint related to the specific GAN architecture that created the Deepfake image, in this paper an extension of our previous work [5] is presented. The fingerprint extraction method based on the Expectation-Maximization Algorithm will be furtherly discussed focusing on its capabilities to extract the Convolutional Traces (CT) embedded by the generative process. The CT could be employed in many related classification tasks but in this paper the finalized pipeline for fakeness detection is finalized with the adoption of a Random Forest classifier. Moreover, the method was deeply tested for robustness with many attacks carried out on images before the extraction of the CT. Also generalizing was demonstrated by testing real images against images generated by ten different GAN architectures, which is the widest test carried out on the task till today. Comparison with state-of-the-art methods demonstrated that the overall approach achieves in almost all cases best classification results. Moreover, we would like to highlight that different state-of-the-art methods for Deepfake detection used approaches based on CNN and these are extremely computationally demanding (both for hardware and for time needed), while the proposed approach achieves

excellent classification results using only the CPU power of a common laptop.

The remainder of this paper is organized as follows: Section II presents state-of-the-art Deepfake generation and detection methods. The proposed approach to extract the Convolutional Trace is described in Section III. Classification phase and experimental results are reported in Section IV. In Section V the proposed approach is demonstrated to be robust to different attacks. Finally, obtained classification results were compared with recent state-of-the-art methods in Section VI. Section VII concludes the paper.

II. RELATED WORKS

Deepfakes are generally created by techniques based on Generative Adversarial Networks (GANs) firstly introduced by Goodfellow *et al.* [3]. In [3], authors proposed a new framework in which two models simultaneously train: a generative model G , that captures the data distribution, and a discriminative model D , able to estimate the probability that a sample comes from the training data rather than from G . The training procedure for G is to maximize the probability of D making a mistake thus resulting in a min-max two-player game. Mathematically, the generator accepts a random input z with density p_z and returns an output $x = G(z, \Theta_g)$ according to a certain probability distribution p_g (Θ_g represents the parameters of the generative model). The discriminator, $D(x, \Theta_d)$ computes the probability that x comes from the distribution of training data p_{data} (Θ_d represents the parameters of the discriminative model). The overall objective is to obtain a generator, after the training phase, which is a good estimator of p_{data} . When this happens, the discriminator is “deceived” and will no longer be able to distinguish the samples from p_{data} and p_g ; therefore p_g will follow the targeted probability distribution, i.e. p_{data} . Figure 2 shows a simplified description of a GAN framework. In the case of Deepfakes, G can be thought as a team of counterfeiters trying to produce fake currency, while D stands to the police, trying to detect the malicious activity. G and D can be implemented as any kind of

generative model, in particular when deep neural networks are employed results become extremely accurate. Through recent years, many GAN architectures were proposed for different applications e.g., image to image translation [6], image super resolution [7], image completion [8], and text-to-image generation [9].

A. DEEPPAKE GENERATION TECHNIQUES FOR FACES

Advances in GAN architectures lead to different works dealing with human faces. STARGAN, created by Choi *et al.* [10], is a method capable of performing image-to-image translations on multiple domains using a single model (e.g, change hair color, facial expression). Many methods work in the latent space representation in order to set constraints to the attributes to be modified, an example is ATTGAN, created by He *et al.* [11]. Cho *et al.* [12] proposed the “group-wise deep whitening-and coloring method” (GDWCT) for a better styling capacity, obtaining a great improvement in the image translation and style transfer task in terms of computational efficiency and quality of generated images. The stage changes when surprising results of Deepfake images were obtained by Style Generative Adversarial Network (STYLEGAN) [13]. STYLEGAN was used to create the so-called “this person does not exist” website.⁶ Moreover, a few imperfect artifact created by STYLEGAN were fixed by Karras *et al.* [14] with improvements to the generator (including re-designed normalization, multi-resolution, and regularization methods), creating the even more realistic images with the so called STYLEGAN2.

B. DEEPPAKE DETECTION METHODS

A starting point to detect Deepfakes is indeed the analysis in the Fourier domain which is a well known technique to find anomalies for image forensics experts [15]. Indeed, some Deepfake images, in the Fourier domain, after being processed by a Discrete Fast Fourier Transform, show abnormal frequencies distributions. This preliminary insight was detected by Guarnera *et al.* [4] in which the authors tried to roughly detect Deepfakes by means of well-known forgery detection tools ([2], [15], [16]) with only few insights for future works as results. The analysis in the Fourier domain was employed by Zhang *et al.* [17] in a rather naive strategy which delivered in any case good performances. Later, an interesting work known as FakeSpotter was proposed by Wang *et al.* [18]. They described a new method based on monitoring neuron behaviors of a dedicated CNN to detect faces generated by Deepfake technologies. The comparison with Zhang *et al.* [17] demonstrated an average detection accuracy of more than 90%

Wang *et al.* [19] trained a ResNet-50 to discriminate real images from those generated by ProGAN [20] and demonstrated that the trained model is able to generalize for the detection of Deepfakes generated by other architectures than ProGAN. They also demonstrate to achieve good

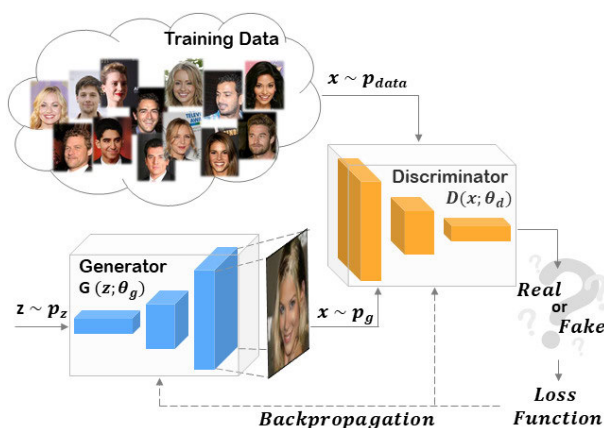


FIGURE 2. Schematic description of a GAN learning framework.

⁶<https://thispersondoesnotexist.com/>

robustness to JPEG compression, spatial blurring and scaling transformations.

Jain *et al.* [21] proposed a work known as DAD-HCNN, a new framework based on a hierarchical classification pipeline composed of three levels to distinguish respectively real Vs altered images (first level), retouched Vs GAN's generated images (second level) and finally, the specific GAN architecture (third level). From the conducted tests, the framework can detect retouching and GANs generated images with high accuracy.

A reference dataset was introduced by Rossler *et al.* [22] as a benchmark for fake detection. It is called FaceForensics++, and is based mainly on four manipulation methods: two computer graphics-based methods (Face2Face [23], FaceSwap⁷) and 2 learning-based approaches (DeepFakes,⁸ NeuralTextures [24]).

By roughly considering literature in the field, it seems like that Deepfake detection is an easy task, already solved. However, analytical techniques based on frequency domain still lack of accuracy and CNN techniques while achieving good results tend to discriminate semantics more than GAN-specific traces. Moreover CNN techniques are computationally intensive and difficult to be understood or controlled [25]. To overcome this, Guarnera *et al.* [5] proposed a new analytical solution to extract a unique fingerprint from images that was demonstrated to be specific to the GAN that generated the image itself. In this paper, the technique will be presented in the mathematical details with furtherly discussion on robustness and generalization, by means of the many carried out experiments: the widest test cases, as today, in the Deepfake detection task will be presented. For this purpose we employed ten of the most famous and effective Deepfake generation architectures: CYCLEGAN [6], STARGAN [10], ATTGAN [11], GDWCT [12], STYLEGAN [13], STYLEGAN2 [14], PROGAN [20], FACEFORENSICS++ [22], IMLE [26] and SPADE [27]. Figure 3 resumes the differences of these techniques in terms of image size, datasets used as input, goal and examples of generated images. For each architecture 2000 images were generated.

III. EXTRACTING CONVOLUTIONAL TRACES

Generative Adversarial Networks (GAN) are used to generate Deepfakes. Once trained, the fundamental element involved in the image creation is the generator G which is composed of Transpose Convolution layers [28]. They apply kernels to the input image, similarly to kernels in Convolutional Layer but they act inversely in order to obtain an output larger but proportional to the input dimensions. Thus, the image creation pipeline is different from the pipeline commonly used in a camera device in which each step introduces typical noise that is then used for naive image forgery detection [15]. However, the image creation process related to the Transpose

Convolution layers of GAN should be consistent and identifiable in local correlations of pixels in the spatial RGB space. To find these traces, an Expectation-Maximization (EM) algorithm [29] was employed in order to define a conceptual mathematical model able to capture the pixel correlation in the images (e.g. spatially) and discriminate between two distributions: the expected one (natural) and others (possibly Deepfake). The result of EM is a feature vector representing the structure of the Transpose Convolution Layers employed during the generation of the image, encoding in some sense if such image is a Deepfake or not, thus it can be called *Convolutional Trace* (CT).

The CT extraction techniques works as follows. The initial goal is to extract a description, from input image I , able to numerically represent the local correlations between each pixel in a neighbourhood. This can be done by means of convolution with a kernel k of $N \times N$ size:

$$I[x, y] = \sum_{s, t = -\alpha}^{\alpha} k_{s, t} * I[x + s, y + t] \quad (1)$$

In Equation 1, the value of the pixel $I[x, y]$ is computed considering a neighborhood of size $N \times N$ of the input data. It is clear that the new estimated information $I[x, y]$ mainly depends on the kernel used in the convolution operation, which establishes a mathematical relationship between the pixels. For this reason, our goal is to define a vector k of size $N \times N$ able to capture this hidden and implicit relationship which characterizes the forensic trace we want to exploit.

Let's assume that the element $I[x, y]$ belongs to one of the following models:

- M_1 : when the element $I[x, y]$ satisfies Equation 1;
- M_2 : otherwise.

The EM algorithm is employed with its two different steps:

- 1) **Expectation step**: computes the (density of) probability that each element belongs to model (M_1 or M_2);
- 2) **Maximization step**: estimates the (weighted) parameters based on the probabilities of belonging to instances of (M_1 or M_2).

Let's suppose that M_1 and M_2 have different probability distributions with M_1 Gaussian distribution with zero mean and unknown variance and M_2 uniform. In the Expectation step, the Bayes rule that $I[x, y]$ belongs to the model M_1 is computed as follows:

$$\begin{aligned} & Pr\{I[x, y] \in M_1 \mid I[x, y]\} \\ &= \frac{Pr\{I[x, y] \mid I[x, y] \in M_1\} * Pr\{I[x, y] \in M_1\}}{\sum_{i=1}^2 Pr\{I[x, y] \mid I[x, y] \in M_i\} * Pr\{I[x, y] \in M_i\}} \quad (2) \end{aligned}$$

where the probability distribution of M_1 which represents the probability of observing a sample $I[x, y]$, knowing that it was generated by the model M_1 is:

$$Pr\{I[x, y] \mid I[x, y] \in M_1\} = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(R[x, y])^2}{2\sigma^2}} \quad (3)$$

where

$$R[x, y] = \left| I[x, y] - \sum_{s, t = -\alpha}^{\alpha} k_{s, t} I[x + s, y + t] \right|. \quad (4)$$

⁷<https://github.com/MarekKowalski/FaceSwap/>

⁸<https://github.com/deepfakes/faceswap/>











<p>STARGAN Image-to-image translations on multiple domains using one model</p>	<p>ATTGAN Transfers face attributes with constraints</p>
<p>Input: CELEBA Output image size: 256x256 #Images Generated: 2000</p>	<p>Input: CELEBA Output image size: 256x256 #Images Generated: 2000</p>
	
<p>GDWCT Improves the styling capability</p>	<p>IMLE Synthesizes images given a semantic layout</p>
<p>Input: CELEBA Output image size: 216x216 #Images Generated: 2000</p>	<p>Input: GTA Output image size: 512x216 #Images Generated: 2000</p>
	
<p>STYLEGAN Transfers semantic content from a source domain to a target domain characterized by a different style</p>	<p>FACE-FORENSICS++ It is not an architecture but a dataset of manipulated videos with four methods</p>
<p>Input: FFHQ Output image size: 1024x1024 #Images Generated: 2000</p>	<p>Input: Youtube Videos Output image size: min: 162x162 max: 895x895 #Images Generated: 2000</p>
	
<p>STYLEGAN2 Improves STYLEGAN quality with the same task</p>	<p>PROGAN Creates images starting from low quality details</p>
<p>Input: FFHQ Output image size: 1024x1024 #Images Generated: 2000</p>	<p>Input: LSUN Output image size: 256x256 #Images Generated: 2000</p>
	
<p>CYCLEGAN Image-to-image translation for everything</p>	<p>SPADE Synthesizes photorealistic images from a semantic layout</p>
<p>Input: Cityscapes, CMP Facade, Google Maps, Zappos50K, ImageNet, Flickr API. Output image size: 256x256 #Images Generated: 2000</p>	<p>Input: ADE20K Output image size: 256x256 #Images Generated: 2000</p>
	

FIGURE 3. Details for each image set used in this paper. On the right of each deep architecture’s name is reported a brief description. *Input* represents the dataset used for both training and test phase of the respective architecture. *Image size* describes the image size of the generated Deepfakes dataset. As regards FACEFORENSICS++ is concerned that for each video frame, the patch referring to the face, is detected and extracted automatically. This patch could have different sizes. *#Images Generated* describes the total number of images taken into account for the considered architecture. Finally, image examples are reported.

The variance value σ^2 , which is still unknown, is then estimated in the Maximization step. Once defined if $I[x, y]$

belongs to model M_1 (or M_2), the values of the vector \bar{k} are estimated using Least Squares method, minimizing

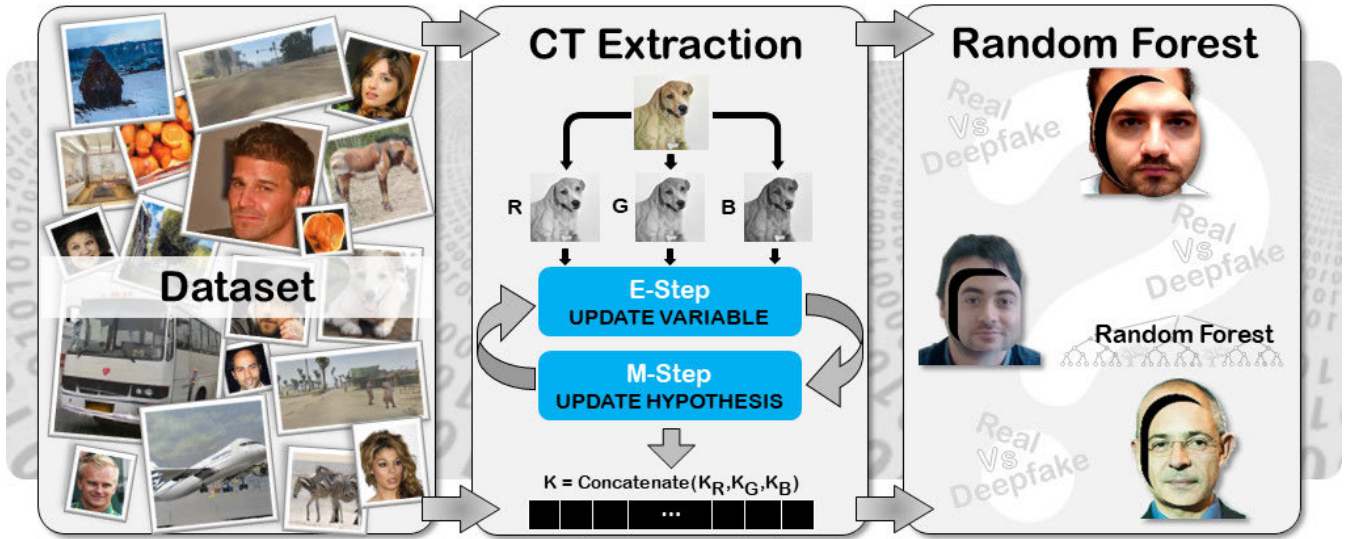


FIGURE 4. Overall finalized Deepfake detection pipeline. The dataset block represents an overview of input data used in this work (Real and Deepfake images). For each image we apply EM algorithm on every channel (R,G,B) obtaining K_R, K_G, K_B feature vectors; the concatenation of them gives the final representation (K) of the input image: the so called Convolutional Trace (CT). Finally, the CT is employed to discriminate real from Deepfake images by means of Random Forest.

the following:

$$E(\vec{k}) = \sum_{x,y} w[x,y] \left(I[x,y] - \sum_{s,t=-\alpha}^{\alpha} k_{s,t} I[x+s,y+t] \right)^2 \quad (5)$$

where $w \equiv Pr\{I[x,y] \in M_1 \mid I[x,y]\}$ (2). This error function (5) is minimized by computing the gradient of vector \vec{k} . The update of $k_{i,j}$ is carried out by computing the partial derivative of (5) as follows:

$$\frac{\partial E}{\partial k_{i,j}} = 0 \quad (6)$$

Hence, the following linear equations system is obtained:

$$\sum_{s,t=-\alpha}^{\alpha} k_{s,t} \left(\sum_{x,y} w[x,y] I[x+i,y+j] I[x+s,y+t] \right) = \sum_{x,y} w[x,y] I[x+i,y+j] I[x,y] \quad (7)$$

The two steps of the EM algorithm are iteratively repeated. The algorithm is applied to each channel of the input image (RGB color space).

The obtained feature vector \vec{k} is the desired CT and has dimensions dependent on parameter α . Note that the element $k_{0,0}$ will always be set equal to 0 ($k_{0,0} = 0$). Thus, for example, if a kernel k with 3×3 size is employed, the resulting \vec{k} will be a vector of 24 elements (since the values $k_{0,0}$ are excluded). This is obtained by concatenating the features extracted from each of the three RGB channels.

The computational complexity of the EM algorithm can be estimated to be linear in d (the number of characteristics of the input data taken into consideration), n (the number of objects) and t (the number of iterations) making it easily to be computed in seconds on a common laptop.

Two aspects are of extreme importance: (i) the proposed CT extraction technique does not need training, it is applied on images and extracts a discriminative feature vector; (ii) the CT extraction is not a deep learning architecture, thus it is not able to encode high level information such semantics. This will be demonstrated in the following Sections with experimental tests.

IV. CLASSIFICATION OF DEEPAKES

In this Section, the Convolutional Trace (CT) extracted by means of the technique presented in Section III, will be demonstrated to have great discriminative power for the Deepfake detection task. Moreover, the independence on image semantics will be demonstrated in this Section by testing against Deepfakes not representing merely faces.

Experiments were carried out considering images created by STARGAN [10], ATTGAN [11], GDWCT [12], STYLEGAN [13], STYLEGAN2 [14] and FACEFORENSICS++ [22] for Deepfake of faces in conjunction with other four Deepfake architectures not dealing with faces: CYCLEGAN [6], PROGAN [20], IMLE [26] and SPADE [27]. Figure 3 shows a brief presentation of the employed images, the techniques, targets, semantics, etc. by reporting also details about training and testing purposes. All images employed in this study are available at <https://iplab.dmi.unict.it/mfs/FightingDeepfake>

STYLEGAN images⁹ and STYLEGAN2 images¹⁰ were downloaded from the official websites, while, for images of the other architectures, the pre-trained models were employed to generate them. The CT was extracted from all the images

⁹<https://drive.google.com/drive/folders/STYLEGAN>

¹⁰<https://drive.google.com/drive/folders/STYLEGAN2>

TABLE 1. Overall accuracy between CELEBA vs. each of the considered GAN. Results are presented w.r.t. all the different kernel sizes (3×3 , 5×5 , 7×7) and with different classifiers: KNN, with $k \in \{3, 5, 7, 9, 11, 13\}$; Linear SVM, Linear Discriminant Analysis (LDA).

	ATTGAN			CYCLEGAN			FACEFORENSICS++			GDWCT			IMLE		
	Kernel Size			Kernel Size			Kernel Size			Kernel Size			Kernel Size		
	3x3	5x5	7x7	3x3	5x5	7x7	3x3	5x5	7x7	3x3	5x5	7x7	3x3	5x5	7x7
3-NN	92.62	82.92	81.51	93.59	86.74	87.29	97.31	84.74	80.70	91.38	72.19	72.49	97.76	97.19	94.73
5-NN	92.99	84.47	80.21	93.33	86.64	87.63	96.98	84.21	78.57	91.08	75.03	73.54	97.69	96.93	94.64
7-NN	92.99	85.20	80.47	93.25	85.85	86.50	96.56	83.68	79.07	91.08	75.60	75.53	97.39	96.85	94.64
9-NN	92.71	84.68	80.47	93.25	85.66	86.84	96.56	82.95	79.07	91.58	75.71	75.00	97.24	96.59	94.25
11-NN	92.90	84.68	79.56	93.00	85.07	86.50	96.47	81.89	79.07	91.48	75.48	74.34	96.94	96.50	94.06
13-NN	92.52	84.99	78.91	92.32	84.97	85.94	96.31	81.79	79.07	91.28	75.94	73.68	96.42	96.59	94.06
SVM	90.19	88.51	87.11	78.90	83.20	90.33	92.70	84.74	83.33	89.30	77.30	80.95	95.67	97.27	95.40
LDA	90.47	87.47	86.59	77.05	83.30	89.65	92.28	84.53	81.70	88.90	77.41	82.01	95.00	96.59	94.64

	PROGAN			SPADE			STARGAN			STYLEGAN			STYLEGAN2		
	Kernel Size			Kernel Size			Kernel Size			Kernel Size			Kernel Size		
	3x3	5x5	7x7	3x3	5x5	7x7	3x3	5x5	7x7	3x3	5x5	7x7	3x3	5x5	7x7
3-NN	95.70	83.38	78.76	96.72	78.35	84.96	88.40	82.08	83.73	94.62	99.48	98.95	96.58	98.06	98.82
5-NN	95.85	82.24	81.08	96.64	78.35	85.15	88.10	82.31	83.46	95.29	99.35	99.12	96.91	98.06	99.15
7-NN	95.54	83.86	82.08	96.27	79.20	85.15	87.88	81.63	82.41	94.72	99.35	99.12	96.80	97.93	99.32
9-NN	95.47	83.10	82.19	95.90	80.05	85.06	88.47	82.42	82.28	94.52	99.35	99.12	96.58	97.93	99.15
11-NN	95.16	82.43	81.53	95.90	79.11	83.81	88.54	82.42	82.28	94.24	99.35	99.12	96.69	97.93	99.15
13-NN	95.39	83.10	82.19	95.90	79.71	84.20	88.25	82.08	82.94	94.14	99.35	99.12	96.48	97.93	99.15
SVM	86.78	80.72	85.18	90.00	83.63	89.46	88.54	84.43	90.55	93.56	99.22	98.77	96.26	99.64	99.32
LDA	86.47	80.91	83.96	88.58	82.69	88.70	87.80	84.32	89.76	93.18	98.82	99.30	96.69	99.03	98.82

with kernels of increasing sizes (3, 5 and 7). The CT obtained was employed as input feature vector for different naive classifiers (K-NN, SVM, LDA) with different tasks: (i) discriminating an authentic image from one generated by a specific GAN and (ii) discriminating authentic images from Deepfakes (binary classification - Real Vs Deepfake images generated by all the 10 techniques). We achieved the best classification solution by employing Random Forest as a final binary classifier, thus finalizing the pipeline (Figure 4).

Let's first analyse the discriminative power of the CT in order to distinguish authentic images from each of the considered GAN. Figure 5 shows a visual representation by means of t-SNE [30]: it is possible to notice how Deepfakes can be "linearly" separable from authentic samples. Moreover, in most cases the separation is utterly clear. Figure 5 visually demonstrates the discriminative power of the extracted CT which, if used as feature vector in a classification task, obtains excellent results as expected. All the classification results are reported in Table 1. In particular, it is possible to note that:

- **CELEBA Vs ATTGAN**: the maximum classification accuracy of 92.99%, was obtained with KNN (with $K = 5, 7$), and kernel size of 3×3 .
- **CELEBA Vs CYCLEGAN**: the maximum classification accuracy of 93.59%, was obtained with KNN (with $K = 3$), and kernel size of 3×3 .
- **CELEBA Vs FACEFORENSICS++**: the maximum classification accuracy of 97.31%, was obtained with KNN (with $K = 3$), and kernel size of 3×3 .
- **CELEBA Vs GDWCT**: the maximum classification accuracy of 91.58%, was obtained with KNN (with $K = 9$) and kernel size of 3×3 .
- **CELEBA Vs IMLE**: the maximum classification accuracy of 97.76%, was obtained with KNN (with $K = 3$) and kernel size of 3×3 .

- **CELEBA Vs PROGAN**: the maximum classification accuracy of 95.85%, was obtained with KNN (with $K = 5$) and kernel size of 3×3 .
- **CELEBA Vs SPADE**: the maximum classification accuracy of 96.72%, was obtained with KNN (with $K = 3$) and kernel size of 3×3 .
- **CELEBA Vs STARGAN**: the maximum classification accuracy of 90.55%, was obtained with linear SVM, and kernel size of 7×7 .
- **CELEBA Vs STYLEGAN**: the maximum classification accuracy of 99.48%, was obtained with KNN - $K = 3$, and kernel size of 5×5 .
- **CELEBA Vs STYLEGAN2**: the maximum classification accuracy of 99.64%, was obtained with linear SVM, and kernel size of 5×5 .

This leads to an empirical hypothesis: the kernel size used by output layers in Deepfake generation techniques is related to the kernel size parameter employed by the CT extraction approach. However, it has to be noted that, on average the kernel size of 3×3 achieves best results among all the classification tests.

Another interesting insight is that the extracted CT is able to discriminate between images from STYLEGAN and STYLEGAN2: a binary test carried out to discriminate between images from the two "similar" techniques achieved a maximum accuracy of 99.31% (Table 2). As stated by the authors of the STYLEGAN2 architecture, they have only updated parts of the generator G , in order to remove imperfections of the original STYLEGAN. This further confirms the former hypothesis, since even a slight modification of G , leaves different traces in the images generated and the CT is able to extract such fingerprint.

We also employed binary classification between real images and Deepfakes coming from all the 10 architectures

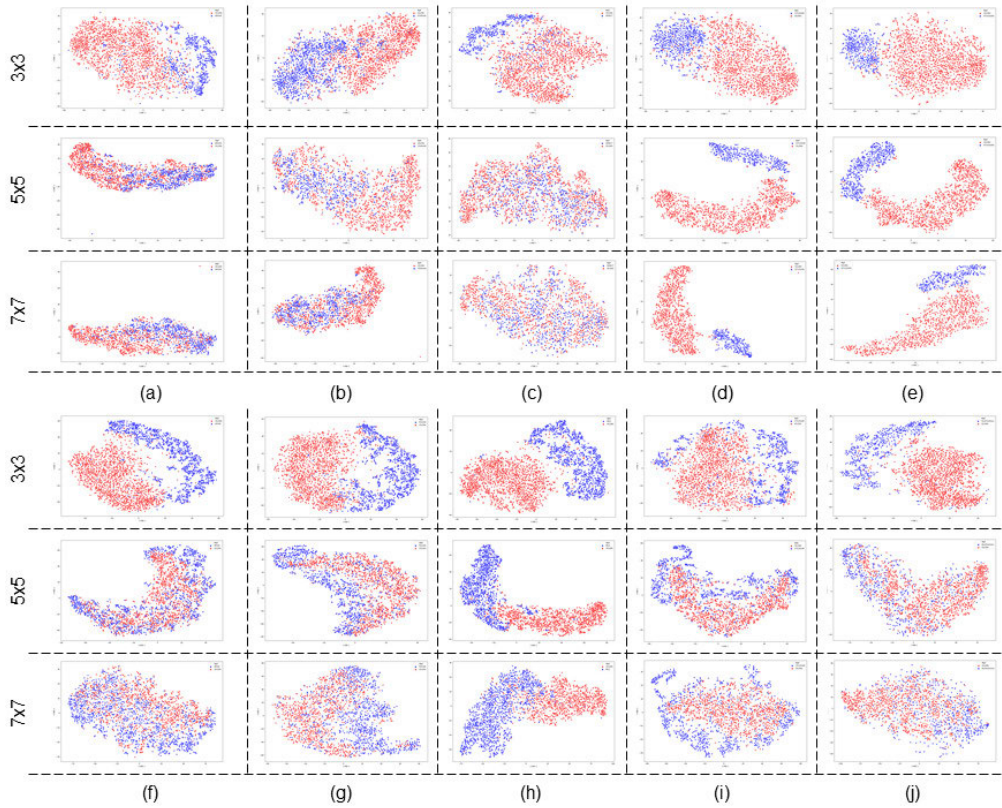


FIGURE 5. Two-dimensional t-SNE representations (CELEBA: red; DeepNetwork: blue) of all kernel sizes for each classification task: (a) CELEBA – ATGAN; (b) CELEBA – STARGAN; (c) CELEBA – GDWCT; (d) CELEBA – STYLEGAN; (e) CELEBA – STYLEGAN2; (f) CELEBA - SPADE; (g) CELEBA - PROGAN; (h) CELEBA - IMLE; (i) CELEBA - CYCLEGAN; (j) CELEBA - FACEFORENSICS++.

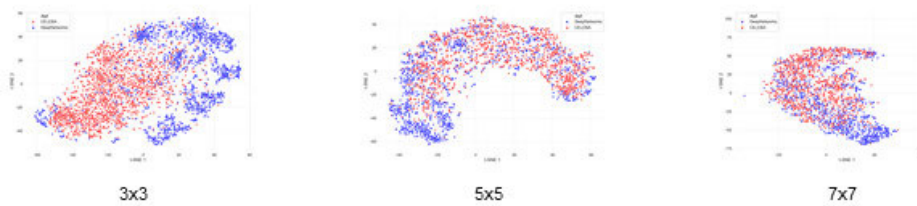


FIGURE 6. Two-dimensional t-SNE representation (CELEBA: red; All 10 DeepNetworks: blue) of a binary classification problem (with different kernel size): CELEBA Vs All 10 DeepNetworks.

TABLE 2. Accuracy values for binary test between STYLEGAN and STYLEGAN2 with different classifiers and kernel sizes (3 × 3, 5 × 5, 7 × 7).

	STYLEGAN Vs STYLEGAN2		
	Kernel Size		
	3x3	5x5	7x7
3-NN	89.36	90.51	87.24
5-NN	89.56	89.87	85.52
7-NN	89.16	90.93	87.59
9-NN	88.55	89.87	87.93
11-NN	88.35	90.30	87.24
13-NN	89.36	89.66	87.93
SVM	91.77	99.16	99.31
LDA	91.16	98.73	98.28

taken into account. At first, another t-SNE representation was built in order to understand sample separability and

distribution in two-dimensional plane. Figure 6 shows that, in this case, samples cannot be linearly separated thus we carried out tests looking for non-linear classifiers. Indeed, final results demonstrated and confirmed such insights. Best accuracy score was obtained by employing Random Forest properly chosen as the final step of the Deepfake detection pipeline (Figure 4) with a solid 98% of accuracy (Table 3) obtained in our tests.

In this Section, experimental results and t-SNE visualizations demonstrated the discriminative power of the CT extracted from Deepfakes. Moreover, the CT achieves good results in detecting Deepfakes not representing faces, hence demonstrating CT being independent to semantics. To further evaluate the proposed pipeline we employed an additional classification test: detecting Deepfakes created by the famous mobile app FACEAPP.

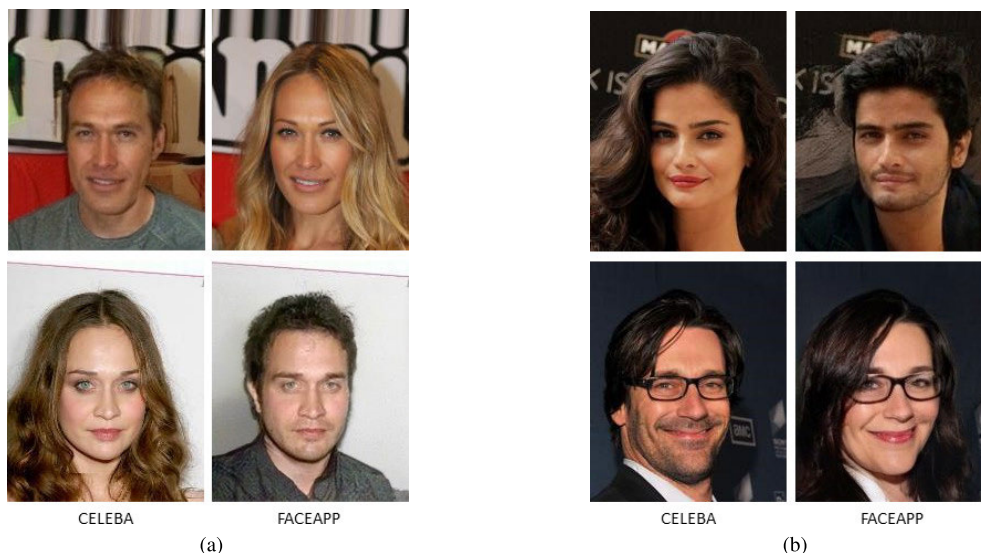


FIGURE 7. (a) Example of images generated by FACEAPP with correct classification (deepfake) (b) Example of images generated by FACEAPP with incorrect classification (real).

TABLE 3. Accuracy values obtained in the binary classification task between Real images vs. images generated by 10 Deepfake architectures. Results are reported with different kernel sizes (3×3 , 5×5 , 7×7) and classifiers trained on 70% of the dataset and tested on the remaining part. Results are the average accuracy value obtained on a 5-fold cross validation test.

	CELEBA Vs DeepNetworks		
	Kernel Size		
	3x3	5x5	7x7
3-NN	89.80	77.38	78.63
5-NN	90.79	77.20	77.80
7-NN	90.44	76.47	78.39
9-NN	90.30	77.20	78.28
11-NN	89.80	77.29	77.45
13-NN	89.73	77.66	77.69
SVMLinear	84.14	76.28	80.28
SVM Sigmoid	58.57	61.36	63.52
SVMrbf	91.22	80.04	80.87
SVM Poly	88.74	78.66	78.87
LDA	83.50	77.38	78.98
Random Forest	98.07	93.81	91.22

Recently, the mobile application called FACEAPP is having a lot of success due to the ability to change features of the input image of a face such as gender, age, hair style, etc. The images thus produced are utterly realistic. Hence, a test for automatic detection of Deepfakes produced by FACEAPP has been carried out employing the CT extraction method and the Random Forest classifier already trained for the test previously described. No further training was done on FACEAPP images. For experiments a dataset of Deepfake images was created starting from CELEBA images by using the Android version of FACEAPP (we employed the paid version that does not introduce watermarks on images): 471 images were generated with FACEAPP by applying gender swap on original images. CTs were extracted with kernel size 3×3 and employed as input for the pre-trained Random Forest classifier. Among the

471 images, 437 were correctly classified as Deepfakes while 34 images were classified as real faces. Figure 7a shows two examples of correct classifications while Figure 7b shows two examples of misclassification. It has to be noted that incorrect classifications are probably due to low light conditions or too few changes in the original images thus making difficult to extract a discriminative CT. According to the reported results we proved the effectiveness of the proposed Deepfake detection technique in a real-case scenario.

V. ROBUSTNESS EXPERIMENTS

Finally, we introduce further tests about overall robustness. A series of attacks were made at different Deepfake images of faces generated by ATTGAN, GDWCT, STARGAN, STYLEGAN and STYLEGAN2 and real images (CELEBA). In particular, the following attacks were carried out:

- 1) *Adding one rectangle with different sizes, positions and colors at random*: in this way details are removed. Since the CT extracts information from pixel correlations, the addition of this rectangle could lead to errors. This could happen specifically for STARGAN or ATTGAN considering that they change only few elements in a face (e.g. hair color) and if these elements are removed by the rectangle low classification accuracy values are expected;
- 2) *Adding Gaussian Blur with different kernel sizes (3×3 , 9×9 , 15×15)*: the noise added to the images could destroy the pixels correlation created by Deepfake architectures and remove the CT;
- 3) *Rotating images by 45, 90, 180 degrees*: rotations could lead to interpolation transformation with modification on CTs similar to the Gaussian blur attack;
- 4) *Scaling images (+50%, -50%)*: due to the interpolation operations carried out, information will be added or removed. CT extracted from images with high details

		CELEBA	ATTGAN	GDWCT	STARGAN	STYLEGAN	STYLEGAN2
Random Square							
Gaussian Blur	Kernel Size 3x3						
	Kernel Size 9x9						
	Kernel Size 15x15						
Rotation	45°						
	90°						
	180°						
Scaling	+50%						
	-50%						
JPEG Compression	Quality Factor = 50						

FIGURE 8. Examples of real (CELEBA) and deepfake images of faces (ATTGAN, GDWCT, STARGAN, STYLEGAN, STYLEGAN2) with six different kind of attacks: *Random Square*, *Gaussian Blur*, *Rotation*, *Scaling* and *JPEG Compression*.

TABLE 4. Robustness to attacks: table reports accuracy values obtained (percentage) for the binary classification task (real vs. Deepfakes) employing the final classification solution for each different kernel size (3×3 , 5×5 , 7×7). The final classifier was trained on the augmented dataset (70% of data for the training set) and 5-fold cross-validated. The first row represents the accuracy obtained by the trained *robust* classifier without any attack.

	ATTGAN			GDWCT			STARGAN			STYLEGAN			STYLEGAN2		
	Kernel Size			Kernel Size			Kernel Size			Kernel Size			Kernel Size		
	3x3	5x5	7x7	3x3	5x5	7x7	3x3	5x5	7x7	3x3	5x5	7x7	3x3	5x5	7x7
Raw Images	92.99	88.51	87.11	91.58	77.41	82.01	88.54	84.43	90.55	95.29	99.48	99.30	96.91	99.64	99.32
Random Square	82.54	75.47	75	62.03	61.54	63.27	81.16	78.95	76.19	97.26	100	97.37	99.02	100	100
Gaussian Blur. kernel size = 3x3	77.78	73.58	72.22	56.96	59.38	61.22	73.91	80.7	61.9	93.15	98.33	92.11	96.08	98.81	96.08
Gaussian Blur. kernel size = 9x9	76.19	76.92	68.57	56.96	67.19	61.22	72.46	77.19	64.29	97.26	100	94.59	96.08	97.62	94.12
Gaussian Blur. kernel size = 15x15	80.95	76.92	77.14	64.56	67.69	57.14	82.61	80.7	75.61	97.26	98.33	94.59	100	97.59	98.04
Rotation 45°	90	84.31	85.29	67.53	73.02	66.67	85.29	82.14	87.8	89.04	91.67	91.89	97.4	94.2	97.62
Rotation 90°	100	94.23	100	93.59	92.19	93.75	92.75	92.98	97.56	100	100	97.3	100	100	100
Rotation 180°	83.87	86.54	82.86	74.36	67.19	59.18	84.06	91.23	78.57	100	100	91.89	97.03	98.8	98.04
Scaling +50%	88.71	78.43	91.18	78.21	71.88	68.09	89.71	83.93	90	97.22	100	97.3	99	98.78	100
Scaling -50%	75.81	78.85	77.78	71.79	57.81	68.09	79.71	64.91	64.29	95.83	96.67	100	99.01	97.59	94.23
JPEG Compression	86.69	91.67	91.18	85.17	89.33	84.66	89.17	92.69	92.01	99.5	99.33	97.57	99.49	98.96	98.55

(such as those of STYLEGAN and STYLEGAN2) would be more robust to this type of operation;

- 5) *JPEG compression with quality factor equal to 50*: in general, a compression operation (such as JPEG) removes high frequency information which could be of major importance for the CT discriminative power. Moreover, a JPEG compression with Quality Factor 50 is similar to those applied by social networks such as Facebook or Instant Messengers like Whatsapp [31], making this test another real-case scenario.

Once the above mentioned filters are applied individually to images, the CT extraction method was applied and Real Vs Deepfake classifications carried out against each GAN (e.g. CELEBA_{RandomSquare} Vs GDWCT_{RandomSquare}, CELEBA_{GaussianBlur} Vs STARGAN_{GaussianBlur}, etc.).¹¹

The classification results are reported in Table 4.¹¹ Figure 8 shows an example of images obtained after operations listed before. It is possible to observe that the dataset plays a fundamental role: the output of ATTGAN, STARGAN and GDWCT and the output of STYLEGAN and STYLEGAN2 after the Gaussian Blur operation: images from ATTGAN, STARGAN and GDWC show a greater visible blur (and therefore a worse visual quality with greater lack of details) respect to STYLEGAN and STYLEGAN2 images. This is mainly determined by the capability of STYLEGAN and STYLEGAN2 to create images of a bigger size.

Results reported in Table 4 show that the CT extracted is robust to almost all considered attacks.

¹¹We report in this table the maximum accuracy classification value obtained through k-NN (with $k = \{3, 5, 7, 9, 11, 13\}$), LDA (Linear Discriminant Analysis), SVM (Support Vector Machine) with linear kernel and Random Forest

In particular, as stated before, STYLEGAN and STYLEGAN2 images obtained the best classification accuracy values (Real Vs Deepfake) due to their bigger original size. GDWCT, which creates the smallest images (Figure 3), is the least robust to attacks and maintains a proper accuracy result comparable with results without attacks only for JPEG compression.

However, another interesting insight comes from the rotation attacks: a rotation of 90 degrees anticlockwise, which is a rotation that does not introduce interpolation, unexpectedly produces better classification results for each of the considered Deepfake architecture. This could be related to a specific *major* direction of the CT and should be better investigated in future works.

VI. COMPARISONS WITH DEEPAKE DETECTION METHODS

Section II presented a detailed discussion of the state-of-the-art in the field of Deepfakes and specifically in Section II-B the detection methods available as today were discussed.

While analytical techniques based on frequency domain still lack of accuracy, CNN based techniques seems to achieve good results but tend to be context-dependent, prone to overfitting and provably depending to high-level semantics extracted from images. Moreover, CNN techniques are computationally intensive and difficult to be explained or controlled. In [25] the authors discussed this limit about CNN. We carried out tests with a deep neural network VGG-16¹² - on spatial and frequency domains - to solve the binary classification task (Real Vs All 10 Deepfakes)

¹²<https://github.com/1297rohit/VGG16-In-Keras>

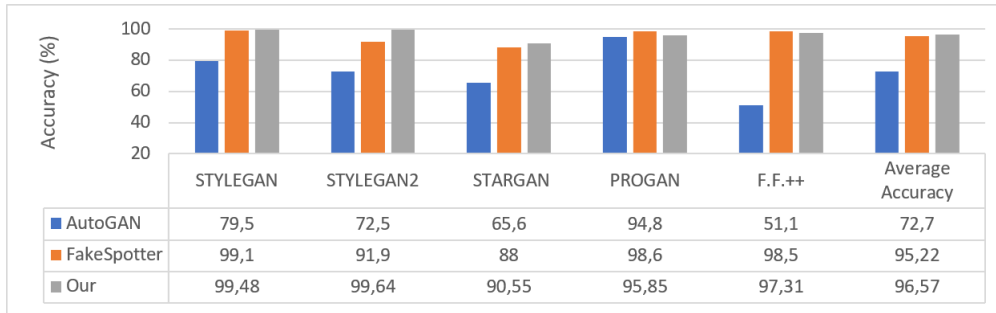


FIGURE 9. Comparison of the proposed approach (Our) vs. FakeSpotter [18] and AutoGAN [17].

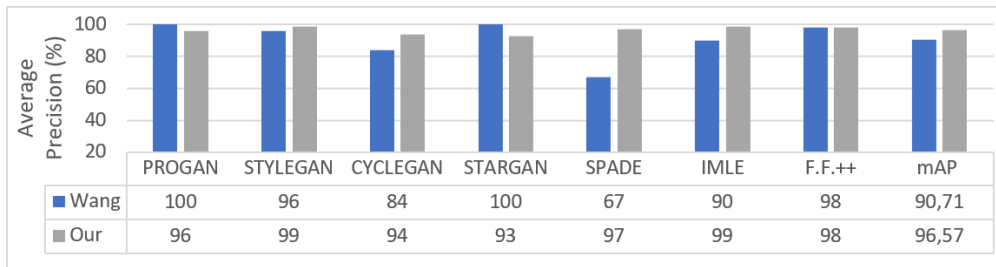


FIGURE 10. Comparison of the proposed approach (Our) vs. Wang et al. [19].

on the datasets described above, obtaining the best result equal to only 53% of accuracy (similar to the random classifier). Better results are achievable by only a complex deep neural network architecture, is what done by recent state-of-the-art methods. The more complex the architecture, the more computing power is required along with the need to understand what high level features the network has used to distinguish Real VS Deepfakes images.

In this Section a detailed discussion is carried out, comparing results obtained by the proposed approach with the best literature methods: Wang et al. [18] (the authors of FakeSpotter), Zhang et al. [17] (the authors of AutoGAN) and Wang et al. [19] were taken into account for comparisons in the Real Vs. Deepfake binary classification task.

For FakeSpotter and AutoGAN, Deepfakes from STYLEGAN, STYLEGAN2, STARGAN, PROGAN and FACEFORENSICS++ architectures were taken into account. Results of this comparison are reported in Figure 9. It is possible to note that, not only we obtained accuracy values of over 90% in all cases, but we overcame FakeSpotter on the average accuracy evaluation. Only in the case of PROGAN and FACEFORENSICS++ we obtained a slightly lower value.

In Wang et al. [19], the following seven Deepfake architectures were taken into account for a fair comparison: PROGAN, STYLEGAN, CYCLEGAN, STARGAN, SPADE, IMLE, FACEFORENSICS++. Wang et al. reported results in the binary classification task as Average Precision between different datasets: images with no data augmentation; images with Gaussian blur added; images JPEG compressed; images

both blurred and JPEG compressed. Figure 10 shows the comparison results obtained by Wang et al. and the proposed approach, reporting the Average Precision (AP) and mean Average Precision (mPA) values for each different Deepfake architecture. It is possible to note that the proposed method obtains better results specifically on Deepfakes of SPADE, IMLE and CYCLEGAN: architectures that do not produce images of faces, furtherly demonstrating the robustness of the extracted CT and classification pipeline to semantics of the image.

VII. CONCLUSIONS AND FUTURE WORKS

In this paper, a finalization of a former work on analysis of Deepfake images was presented. An algorithm based on Expectation-Maximization was employed to extract the Convolutional Traces (CT): a sort of unique fingerprint useable to identify not only if an image is a Deepfake but also the GAN architecture that generated it. The CT extracted is a fingerprint demonstrated to have high discriminative power, robustness to attacks and independence to high-level concepts of images (semantics). Obtained results demonstrate also to overcome the state-of-the-art with a technique simple and fast to be computed. Indeed the CT is related to the generation process of images and further better results can be obtained by rotating input images in order to find the most important direction. This particular hint will be investigated in future works in conjunction with analysis on famous forensics datasets of real images (like DRESDEN, UCID or VISION) which do not focus their contents on human faces.

REFERENCES

- [1] L. Verdoliva, "Media forensics and DeepFakes: An overview," *IEEE J. Sel. Topics Signal Process.*, vol. 14, no. 5, pp. 910–932, Aug. 2020.
- [2] A. Piva, "An overview on image forensics," *Int. Scholarly Res. Notices*, vol. 2013, 2013.
- [3] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
- [4] L. Guarnera, O. Giudice, C. Nastasi, and S. Battiato, "Preliminary forensics analysis of DeepFake images," 2020, *arXiv:2004.12626*. [Online]. Available: <http://arxiv.org/abs/2004.12626>
- [5] L. Guarnera, O. Giudice, and S. Battiato, "DeepFake detection by analyzing convolutional traces," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2020, pp. 666–667.
- [6] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired Image-to-Image translation using cycle-consistent adversarial networks," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2223–2232.
- [7] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi, "Photo-realistic single image super-resolution using a generative adversarial network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 4681–4690.
- [8] S. Iizuka, E. Simo-Serra, and H. Ishikawa, "Globally and locally consistent image completion," *ACM Trans. Graph. (TOG)*, vol. 36, no. 4, pp. 1–14, 2017.
- [9] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee, "Generative adversarial text to image synthesis," 2016, *arXiv:1605.05396*. [Online]. Available: <http://arxiv.org/abs/1605.05396>
- [10] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo, "StarGAN: Unified generative adversarial networks for multi-domain Image-to-Image translation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 8789–8797.
- [11] Z. He, W. Zuo, M. Kan, S. Shan, and X. Chen, "AttGAN: Facial attribute editing by only changing what you want," *IEEE Trans. Image Process.*, vol. 28, no. 11, pp. 5464–5478, Nov. 2019.
- [12] W. Cho, S. Choi, D. K. Park, I. Shin, and J. Choo, "Image-to-image translation via group-wise deep whitening-and-coloring transformation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 10639–10647.
- [13] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 4401–4410.
- [14] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of StyleGAN," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 8110–8119.
- [15] S. Battiato, O. Giudice, and A. Paratore, "Multimedia forensics: Discovering the history of multimedia contents," in *Proc. 17th Int. Conf. Comput. Syst. Technol. (CompSysTech)*, 2016, pp. 5–16.
- [16] O. Giudice, F. Guarnera, A. Paratore, and S. Battiato, "1-D DCT domain analysis for JPEG double compression detection," in *Proc. Int. Conf. Image Anal. Process.* Springer, 2019, pp. 716–726.
- [17] X. Zhang, S. Karaman, and S.-F. Chang, "Detecting and simulating artifacts in GAN fake images," in *Proc. IEEE Int. Workshop Inf. Forensics Secur. (WIFS)*, Dec. 2019, pp. 1–6.
- [18] R. Wang, F. Juefei-Xu, L. Ma, X. Xie, Y. Huang, J. Wang, and Y. Liu, "FakeSpotter: A simple yet robust baseline for spotting AI-synthesized fake faces," 2019, *arXiv:1909.06122*. [Online]. Available: <http://arxiv.org/abs/1909.06122>
- [19] S.-Y. Wang, O. Wang, R. Zhang, A. Owens, and A. A. Efros, "CNN-generated images are surprisingly easy to spot... for now," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 1–10.
- [20] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of GANs for improved quality, stability, and variation," 2017, *arXiv:1710.10196*. [Online]. Available: <http://arxiv.org/abs/1710.10196>
- [21] A. Jain, P. Majumdar, R. Singh, and M. Vatsa, "Detecting GANs and retouching based digital alterations via DAD-HCNN," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2020, pp. 672–673.
- [22] A. Rossler, D. Cozzolino, L. Verdoliva, C. Riess, J. Thies, and M. Niessner, "FaceForensics++: Learning to detect manipulated facial images," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 1–11.
- [23] J. Thies, M. Zollhofer, M. Stamminger, C. Theobalt, and M. Niessner, "Face2Face: Real-time face capture and reenactment of RGB videos," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2387–2395.
- [24] J. Thies, M. Zollhöfer, and M. Nießner, "Deferred neural rendering: Image synthesis using neural textures," *ACM Trans. Graph.*, vol. 38, no. 4, pp. 1–12, Jul. 2019.
- [25] N. Hulzebosch, S. Ibrahim, and M. Worring, "Detecting CNN-generated facial images in real-world scenarios," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2020, pp. 642–643.
- [26] K. Li, T. Zhang, and J. Malik, "Diverse image synthesis from semantic layouts via conditional IMLE," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 4220–4229.
- [27] T. Park, M.-Y. Liu, T.-C. Wang, and J.-Y. Zhu, "Semantic image synthesis with spatially-adaptive normalization," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 2337–2346.
- [28] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," 2015, *arXiv:1511.06434*. [Online]. Available: <http://arxiv.org/abs/1511.06434>
- [29] T. K. Moon, "The expectation-maximization algorithm," *IEEE Signal Process. Mag.*, vol. 13, no. 6, pp. 47–60, Nov. 1996.
- [30] L. van der Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, pp. 2579–2605, Nov. 2008.
- [31] O. Giudice, A. Paratore, M. Moltisanti, and S. Battiato, "A classification engine for image ballistics of social data," in *Proc. Int. Conf. Image Anal. Process.* Springer, 2017, pp. 625–636.



LUCA GUARNERA (Student Member, IEEE) received the master's degree (*cum laude*) in computer science from the University of Catania, in 2017. He is currently an Executive Ph.D. Student in computer science with the University of Catania, where he works for the spin off of the University of Catania "iCTLab s.r.l.", a company operating in the field of digital forensics, privacy and security consulting and software development. He joined IPLab in 2015. In 2017 and 2019,

he participated at the Mohamed Bin Zayed International Robotics Challenge (MBZIRC), an international robotics competition. He participated in four editions (2016–2017–2018–2019) of the International Computer Vision Summer School (ICVSS), an edition (2018) of Medical Imaging Summer School (MISS), and an edition (2018) of Summer School on Signal Processing (S3P). His research interests are computer vision, machine learning, multimedia forensics and its related fields with a particular focused on Deepfake phenomenon.



OLIVER GIUDICE received the degree in computer engineering (*summa cum laude*) from the University of Catania, in 2011, and the Ph.D. degree in mathematics and computer science, in 2017, defending a thesis entitled “Digital Forensics Ballistics: Reconstructing the source of an evidence exploiting multimedia data”. From 2011 to 2014, he was involved in various research projects at the University of Catania in collaboration with the Civil and Environmental Engineering

Department and the National Sanitary System. In 2014, he started his job as a Researcher at the IT Department of Banca d’Italia. For various years since 2011, he collaborated with the IPLab working on multimedia forensics topics and being involved in various forensics cases as a Digital Forensics Expert. Since 2016, he is co-founder of “iCTLab s.r.l.”, spin-off of University of Catania, company that works in the field of digital forensics, privacy and security consulting and software development. His research interests include machine learning, computer vision, image coding, urban security, cryptocurrencies, and multimedia forensics.



SEBASTIANO BATTIATO (Senior Member, IEEE) received the degree (*summa cum laude*) in computer science from the University of Catania, in 1995, and the Ph.D. degree in computer science and applied mathematics from the University of Naples, in 1999. From 1999 to 2003, he was the Leader of the “Imaging” Team, STMicroelectronics, Catania. He joined the Department of Mathematics and Computer Science, University of Catania, as an Assistant

Professor, an Associate Professor, and a Full Professor, in 2004, 2011, and 2016, respectively. He has been the Chairman of the Undergraduate Program in Computer Science, from 2012 to 2017, and a Deputy Rector for education (Postgraduates and Ph.D.) from 2013 to 2016. He is currently a Full Professor of computer science with the University of Catania, where he is also the Scientific Coordinator of the Ph.D. Program in Computer Science and a Deputy Rector for Strategic Planning and Information Systems. He is involved in the research and directorship with the Image Processing Laboratory (IPLab). He coordinates several large scale projects funded by national and international funding bodies and private companies. He has edited six books and coauthored about 250 articles in international journals, conference proceedings, and book chapters. He is a co-inventor of about 25 international patents. His current research interests include computer vision, imaging technology, and multimedia forensics. He has been a regular member of numerous international conference committees. He was a recipient of the 2017 PAMI Mark Everingham Prize for the series of annual ICVSS schools and the 2011 Best Associate Editor Award of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY. He has been the Chair of several international events, including INTELLISYS 2020, SIGMAP 2019-2020 ICIAP 2017, VINEPA 2016, ACIVS 2015, VAAM2014-2015-2016, VISAPP2012-2015, IWCV2012, ECCV2012, ICIAP 2011, ACM MiFor 2010-2011, and SPIE EI Digital Photography 2011-2012-2013. He has been a Guest Editor of several special issues published in international journals. He is an Associate Editor of SPIE *Journal of Electronic Imaging* and *IET Image Processing* journal. He is the Director and Co-Founder of the International Computer Vision Summer School (ICVSS).

• • •