

Exposing Face-Swap Images Based on Deep Learning and ELA Detection [†]

Weiguo Zhang and Chenggang Zhao *

College of Computer Science and Technology, Xi'an University of Science and Technology,
Xi'an 710054, China; wgzhang@scut.edu.cn

* Correspondence: zhaocg@mail.com

[†] Presented at the 5th International Electronic Conference on Entropy and Its Applications, 18–30 November 2019; Available online: <https://ecea-5.sciforum.net/>.

Published: 17 November 2019

Abstract: New developments in artificial intelligence (AI) have significantly improved the quality and efficiency in generating fake face images; for example, the face manipulations by DeepFake are so realistic that it is difficult to distinguish their authenticity—either automatically or by humans. In order to enhance the efficiency of distinguishing facial images generated by AI from real facial images, a novel model has been developed based on deep learning and error level analysis (ELA) detection, which is related to entropy and information theory, such as cross-entropy loss function in the final Softmax layer, normalized mutual information in image preprocessing, and some applications of an encoder based on information theory. Due to the limitations of computing resources and production time, the DeepFake algorithm can only generate limited resolutions, resulting in two different image compression ratios between the fake face area as the foreground and the original area as the background, which leaves distinctive artifacts. By using the error level analysis detection method, we can detect the presence or absence of different image compression ratios and then use Convolution neural network (CNN) to detect whether the image is fake. Experiments show that the training efficiency of the CNN model can be significantly improved by using the ELA method. And the detection accuracy rate can reach more than 97% based on CNN architecture of this method. Compared to the state-of-the-art models, the proposed model has the advantages such as fewer layers, shorter training time, and higher efficiency.

Keywords: artificial intelligence; deep learning; DeepFake detection; ELA detection

1. Introduction

Today, with the popularization of smartphones and various face-swap applications, the manipulation of visual content is becoming more and more common, which has become one of the most critical topics in the digital society. Faces are the main focus of visual content manipulation. There are many reasons for this focus. First of all, face reconstruction and tracking is a relatively mature field in computer vision [1], which is the basis of these editing methods. Also, the human faces play a key role in communications because the human face can emphasize and convey certain information in its own ways [2].

The root of the problem comes from the new generation of generative deep neural networks [3], which are capable of synthesizing videos from a large volume of training data with minimum manual editing. The recent appearance of DeepFake [4] greatly reduces the threshold of face forgery techniques. DeepFake replaces the face in an original video with the face of another person using generative adversary networks (GANs) [5]. Because the GAN models were trained using tens

of thousands of images, it is possible to generate realistic faces that can be spliced into the original video in an almost perfect way. Through suitable post-processing, the resulting video can achieve higher authenticity.

In addition to DeepFake technology, Fake2Face [6] and Faceswap [7] are prominent representatives for facial manipulation. Recently, wide-spread consumer-level applications like ZAO have become popular in China. While face swapping based on simple computer graphics or deep learning is run in real time, DeepFakes need to be trained for every pair of videos, which is a time-consuming and resource-demanding task.

Before the emergence of fake video, it was generally believed that videos were reliable and dependable, and video evidence was widely used in multimedia forensics. However, after the prevalence of fake videos, people's psychological security zone became broken. There is widespread concern that once such fake videos are used for court proof, press and publication, political elections, and television and entertainment, it will become difficult to estimate their impact on people's lives. Some people even think that this technology could hinder the development of society. In this case, the detection and identification of such fake videos, whether for digital media forensics or in ordinary people's lives, have become extremely urgent.

In this paper, we describe a novel model based on the deep learning and error level analysis (ELA) detection, which can effectively distinguish facial images generated by AI from real facial ones. Our experiment is based on a characteristic of DeepFake principle: due to the limitations of computing resources and production time, the DeepFake algorithm can only generate limited resolutions, resulting in two different image compression ratios between the fake face area as the foreground and the original area as the background, which leaves distinctive artifacts.

By using the error level analysis (ELA) detection method, our model can capture such artifacts because the entire image should have roughly the same compression level for JPEG formats. However, if a part of the image has been modified, such as by copy and paste or other removal operations, there will be a significant error level between the tampered part and the surrounding part. At this time, ELA images with different error levels can be generated by ELA method, and the tampered part will be displayed with an obvious white color.

By using the ELA detection method, we can detect the presence or absence of the different compression ratios in the image [8]. We will input the generated ELA image of the real face and fake face into the special convolutional neural network model and train a binary classifier to distinguish whether the image is fake.

2. Related Work

2.1. AI-Based Video Synthesis Algorithms

With 3D computer graphics-based methods, it is easy to generate realistic images/video. Recently, the new deep learning algorithms have developed rapidly, especially those based on generative adversary networks (GANs). Goodfellow et al. [7] first proposed the new generative adversary networks (GANs), which usually consist of two networks—a generator and a discriminator. Face2Face, proposed by Thies et al. [6], is an advanced real-time facial reenactment system that can change the facial movements in video streams, such as videos from the movies.

Recently, some facial image synthesis methods based on deep learning techniques have been proposed. Most of these techniques have the problem of low image resolution. Karras et al. [9] used a progressively growing GAN to improve image quality. Their results include high-quality facial synthesis.

2.2. GAN-Generated Image/Video Detection

With the popularity of face-swap applications, detecting GAN-generated images/videos technology has also made some progress. Li et al. [10] observed that DeepFake faces lack realistic eye blinking because the image collected through the Internet typically does not include photos with closed eyes. Therefore, the lack of eye blinking is detected with a CNN model to expose

DeepFake videos. However, this method can be invalidated by purposely adding images with eyes closed in training.

Li et al. [11] used the color difference between GAN-generated images and real images in non-RGB color spaces to classify them. Afchar et al. [12] trained convolution neural network to directly classify real faces and fake faces generated by DeepFake and Face2face [6]. Although it showed promising performance, the overall approach has its drawbacks. In particular, it needs both true and false images as training data, and generating the fake images using the AI-based synthesis algorithms is less efficient than the simple mechanism for training data generation used in our method. Because extracting features directly from the original image, it needs to go through too many training cycles, resulting in low efficiency.

2.3. Image Tampering Detection

As reliable evidence of judicial identification, digital image authentication technology has made a series of achievements in the field of image tampering detection. Previous methods can be classified according to the image features they aim at, such as Components Factor Analysis (CFA) pattern analysis, local noise estimation, double JPEG localization. Bianchi et al. [13] proposed a probability model for estimates Discrete Cosine Transform (DCT) coefficients and quantization factors. Fu et al. [14] determined whether the image has been tampered by estimating quality factor. Ferrara et al. [15] proposed a model to estimate camera filter mode based on the difference of the variance of prediction error between CFA existing areas (authentic areas) and CFA absent areas (tempered areas). After the Gaussian Mixture Model (GMM) classification, the tampered regions can be localized.

3. Methods

In this section, we will describe the method of detecting facial images forgery in detail. First of all, we analyzed the principle of DeepFake generating face and simulate the process of affine transformation generating a fake face. Then the data sets of real face and fake face are processed by ELA method, and the resulting ELA image will highlight some parts of the original image where the error level is higher than the threshold value, that is, the affine transformation introduced artifacts. Finally, a binary classifier is trained by convolutional neural network (CNN) to distinguish whether the image is fake.

3.1. Data Sets Preprocessing

We analyzed the process of generating fake face by DeepFake. The principle of DeepFake is shown as Figure 1. Due to the limitations of computing resources and production time, the DeepFake algorithm can only generate limited resolutions and then perform affine transformation on those generating images, such as scaling, rotation, and shearing, to match and cover the original face (see Figure 1g–h). This will result in two different image compression ratios between the fake face area as the foreground and the original area as the background, which would leave obvious artifacts.

Our purpose here is to detect the artifacts introduced by the affine face wrapping steps in DeepFake production pipeline. On the other hand, due to DeepFake's need to be trained for each pair of videos, which is a time-consuming and resource-demanding task, we did not use the DeepFake algorithm to create negative examples. Instead, we simplified the process of generating negative examples by simulating the process of generating a face using DeepFake (Figure 1).

Specifically, we took the following steps to generate negative examples, as shown in Figure 2: First, we detect faces in the original image, extract face landmarks from each detected face area, and calculate the transform matrix according to the landmarks. Then, we apply Gaussian blur to the adjusted face. According to the inverse of transform matrix, the face is wrapped back to the original angle and cover on the original face.

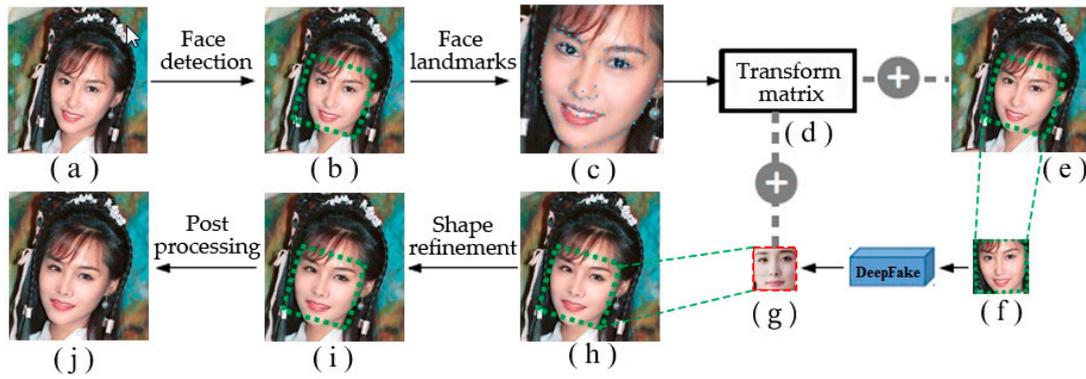


Figure 1. Overview of the DeepFake production pipeline. (a) The source image. (b) The green box is the detected face area. (c) Blue points are the face landmarks. (d) Calculate the transform matrix to wrap face region in (e) to the normalized region (f). (g) The face image generated by the neural network is used to cover the source image (a). (h) Synthesized face wrapped back using the same transform matrix. (i) Post-processing, such as applying boundary smoothing to composite image. (j) The final synthesized image.

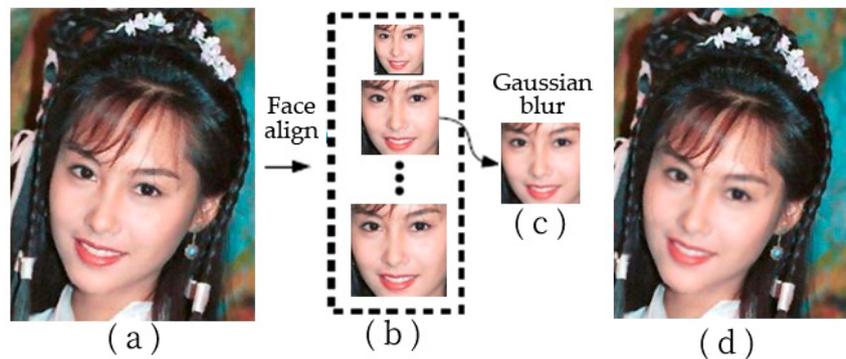


Figure 2. Overview of the process of generating a negative example. (a) The original image. Use dlib to extract the face in (a) and align face with different scales as in (b). We randomly select a scale of face in (b) and apply Gaussian blur as (c), and after the affine transform, cover the original image to generate a negative example.

In order to simulate more different resolutions of face affine transform in reality, we align faces into multiple scales and randomly select one scale to enlarge the training diversity. At the same time, we also use image enhancement technology to simulate different post-processing technology that may exist in the DeepFake process. Our approach also further deals with the shape of the face area affine transformation to cope with the different post-processing techniques.

In addition, we also need to preprocess the images before training. Since the input of convolutional neural network is 128×128 , the size of the image area is not large. Therefore, it is important to retain the most effective and prominent signs of forgery area as our region of interest (RoI). Analyzed as above, there are many trace forgery marks in the surrounding region involved in affine transform, which thereby retains the rectangular region composed of the convex hull of facial landmarks (except the contour of the cheek) and the surrounding area, and removes the remaining part of the image.

Specifically, for all positive and negative examples in the dataset, we only keep the above rectangular region, which is slightly larger than the external rectangular region of the convex landmarks of the face (except the contour of the cheek). We determine the RoIs using face landmarks, as $[y_0 - \hat{y}_0, x_0 - \hat{x}_0, y_1 + \hat{y}_1, x_1 + \hat{x}_1]$, where y_0, x_0, y_1, x_1 indicates the minimum bounding box b which you can cover all the facial landmarks except the cheek contour. The variables $\hat{y}_0, \hat{x}_0, \hat{y}_1, \hat{x}_1$ are random values between $[0, \frac{h}{5}]$ and $[0, \frac{w}{8}]$, where h, w are the height and width of b respectively. The RoIs are resized to 128×128 for the next ELA processing.

3.2. ELA Processing

The error level analysis (ELA) method is one of the techniques for detecting images that have been tampered with. ELA can obtain the compression distortion during lossy image compression. This method detects image tampering by storing images at a specific level of quality and calculating the ratio between compression levels [8]. Typically, this method is performed on images with lossy compression formats, such as JPEG.

When saving images in JPEG format, it will be independent “lossy compression” in units of 8×8 pixels. After lossy compression of a JPEG, there are significant differences between the ELA of the original area and the ELA of the spliced or modified. If the image has not been modified, the compression difference of each 8×8 pixel region is similar. We check the “compression feature” of the tested image with an 8×8 pixel grid. If the image is saved as a whole, the compression feature of the adjacent grid should be an approximately high-frequency white distribution.

On the contrary, if it is saved after editing or modification, the ELA distribution between the grids will have obvious difference characteristics, which is shown as discontinuous high-frequency white distribution. The more times the images are stored or edited, the lower the ELA. The ELA processing effect is shown in Figure 3.



Figure 3. Samples of the original image and tampered images and their error level analysis (ELA) results: The first line are the original image and its ELA image. It can be seen that the compression ratio of the whole image remains the same. The second line is the tampered image and its ELA image. It can be seen that the compression ratio between the tampered face as the foreground and the original image as the background are quite different.

4. Experiments

In this section, we first introduce our dataset and then evaluate our model. In addition, we visualize present our results in order to better understand the proposed model.

4.1. Dataset

Although there are some datasets [16–19] for image tampering detection, they are not suitable for large-scale facial tampering detection because there are not enough tampering samples concentrated in facial areas. The Columbia Image Splicing dataset [16] and Insitute of Automation, Chinese Academy of Sciences (CASIA) [17,18] are large but most of the tampered areas are not human faces. The DSI-1 dataset [19] focuses on facial tampering, but the total number of tampered-with images is only 25. Therefore, it is difficult to train deep learning methods on these datasets to detect facial tampering.

To do this, we used the Milborrow University of Cape Town (MUCT) database, which consists of 3755 facial images and 76 manual facial landmarks. Each compressed file in the data corresponds to a camera, providing more diverse lighting, age, and race than the currently available 2D face database.

We take the 3755 “jpg” format face images in the database as the examples, and the negative examples can be generated by simulating the DeepFake algorithm, as shown in Figure 2, but it requires us to train and run DeepFake, which is a time-consuming and resource-demanding algorithm. Therefore, we use the method in Section 3.1 to generate negative examples dynamically and train them. Dynamic means that instead of generating all the negative examples in advance before the training process, we randomly select half of the positive examples for each training batch and convert them to negative examples according to the process in Figure 2 in order to make the training data more diverse.

4.2. Experiment Setup

For the 128×128 ROI region images generated in the previous step, we use ELA to process them and get their ELA images. The CNN model that we trained uses these ELA images, rather than the original ones. Converting the original image to ELA image is a method to improve the training efficiency of CNN model. Because the ELA image does not contain as much information as the original image, it can improve the efficiency.

The feature generated by ELA image focuses on the part of the original image where the error level is higher than the threshold value. In addition, the pixels in the ELA image are often quite different from the nearby pixels, and even the contrast is very obvious, so the image processed by ELA makes the training CNN model more effective.

Therefore, we train a CNN model to extract the features of the ELA images, then detect whether the input image is forged or not. In the architecture we use, only two convolution layers are required, because the ELA images generated during the conversion process can highlight the characteristics of the original image where the error level is higher than the threshold value. So it is easier to determine whether the image is fake.

The maximum accuracy of the results obtained by our proposed method is 97%. The image of the accuracy curve and the loss function curve can be seen in Figure 4a. The confusion matrix of verification data is shown in Figure 4b.

As shown in Figure 4, our model achieves the best accuracy in the ninth cycle. From the first nine cycles, verify that the value of the loss function starts to be flat and eventually begins to increase, which is a sign of over-fitting. This is also a recognition method of ending training in advance during training, that is, when the verification accuracy value begins to decrease or the verification loss value starts to increase, the training will be stopped.

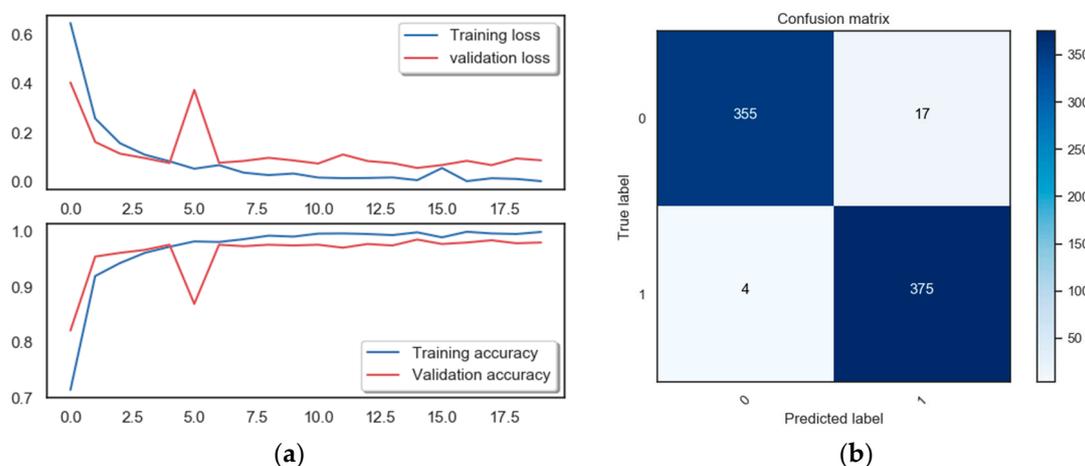


Figure 4. This is the experimental result images. (a) The image of the accuracy curve and the loss function curve; (b) the confusion matrix of verification data.

4.3. Comparison with Other Methods

We compare our method with the method [20] of training directly using CNN without ELA processing. The code for this method is available from the public implementation on GitHub [21]. In this method, the positive and negative examples images are directly input to the network model for training, and this method trained four CNN models—VGG16, ResNet50, ResNet101, and ResNet152. The AUC performance on VGG16, ResNet50, ResNet101, and ResNet152 reached 83.3%, 97.4%, 95.4%, and 93.8%, respectively.

However, compared with our method, this method has the following problems:

1. The deep learning training model cannot explain the deep principle of identification forgery. Our ELA method can explain the principle.
2. This method is too complicated to train. On the one hand, if there is no GPU environment, it will lead to a long training period. On the other hand, it also requires a larger number of samples to participate in the training. Our method using two layers convolution—a MaxPooling layer, and a fully connected layer. An output layer with Softmax can reach 97% accuracy and greatly reduces the training time and the training period.

The advantages of our model are as follows: the number of training periods required to achieve convergence is significantly reduced because the image features processed by ELA make the training more efficient and accelerate the convergence of CNN model. On the other hand, the accuracy of our classification results is very high. This indicates that the features in the image processed by ELA can be successfully used to classify whether the image is fake. Experiments show that the training efficiency of CNN model can be significantly improved by using the ELA method.

5. Conclusions

New developments in AI have significantly improving quality and efficiency in generating false faces. In this work, we studied a new model based on the deep learning that can effectively distinguish facial images generated by AI from real facial images.

We evaluated our method and proved its effectiveness in practice. This indicates that the features in the image processed by ELA can be successfully used to classify whether the image is fake. Experiments show that the training efficiency of a CNN model can be significantly improved by using the ELA method.

As the technology behind DeepFake continues to develop, we will continue to improve detection methods. We want to evaluate and improve the robustness of our detection methods for video compression.

References

1. Zollhöfer, M.; Thies, J.; Garrido, P.; Bradley, D.; Beeler, T.; Pérez, P.; Stamminger, M.; Nießner, M.; Theobalt, C. State of the art on monocular 3d face reconstruction, tracking, and applications. *Comput. Graph. Forum* **2018**, *37*, 523–550.
2. Frith, C. Role of facial expressions in social interactions. *Philos. Trans. R. Soc. B Biol. Sci.* **2009**, *364*, 3453–3458.
3. Liu, M.-Y.; Breuel, T.; Kautz, J. Unsupervised image-to-image translation networks. In Proceedings of the Advances in Neural Information Processing Systems 30 (NIPS 2017), Long Beach, CA, USA, 4–9 December 2017; pp. 700–708.
4. Deepfakes Github. Available online: <https://github.com/deepfakes/faceswap>, note (accessed on 29 October 2018).
5. Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative adversarial nets. In Proceedings of the Advances in Neural Information Processing Systems, Montreal, QC, Canada, 8–13 December 2014; pp. 2672–2680.
6. Thies, J.; Zollhöfer, M.; Stamminger, M.; Theobalt, C.; Nießner, M. Face2Face: Real-Time Face Capture and Reenactment of RGB Videos. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 26 June–1 July 2016; pp. 2387–2395.

7. Faceswap. Available online: <https://github.com/MarekKowalski/FaceSwap/> (accessed:15 February 2018).
8. Jeronimo DC, Borges YCC, Coelho L dos S. Image forgery detection by semi-automatic wavelet soft-Thresholding with error level analysis. *Expert Syst, Appl.* **2017**, *85*, 348–356.
9. Karras, T.; Aila, T.; Laine, S.; Lehtinen, J. Progressive Growing of GANs for Improved Quality, Stability, and Variation. *arXiv* **2017**, arXiv:1710.10196.
10. Li, Y.; Chang, M.C.; Lyu, S. In *ictu oculi: Exposing ai generated fake face videos by detecting eye blinking*. In Proceedings of the IEEE International Workshop on Information Forensics and Security (WIFS), Hong Kong, China, 11–13 December 2018.
11. Li, H.; Li, B.; Tan, S.; Huang, J. Detection of deep network generated images using disparities in, color components. *arXiv* **2018**, arXiv:1808.07276.
12. Afchar, D.; Nozick, V.; Yamagishi, J.; Echizen, I. Mesonet: A compact facial video forgery detection network. In Proceedings of the IEEE International Workshop on Information Forensics and Security (WIFS) 2018, Hong Kong, China, 11–13 December 2018.
13. Bianchi, T.; de Rosa, A.; Piva, A. Improved dct coefficient analysis for forgery localization in jpeg images. In Proceedings of the 2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Prague, Czech Republic, 22–27 May 2011; pp. 2444–2447.
14. Fu, D.; Shi, Y.Q.; Su, W. A generalized Benford’s law for JPEG coefficients and its applications in image forensic. In Proceedings of the Electronic Imaging. International Society for Optics and Photonics, San Jose, CA, USA, 28 January–1 February 2007.
15. Ferrara, P.; Bianchi, T.; de Rosa, A.; Piva, A. Image forgery localization via fine-grained analysis of cfa artifacts. *IEEE Trans. Inf. Forensics Secur.* **2012**, *7*, 1566–1577.
16. Ng, T.-T.; Hsu, J.; Chang, S.-F. *Columbia Image Splicing Detection Evaluation Dataset*; DVMM Laboratory of Columbia University, CalPhotos Digital Library: New York, NY, USA, 2009.
17. Dong, J.; Wang, W.; Tan, T. Casia image tampering detection evaluation database. In Proceedings of the 2013 IEEE China Summit and International Conference on Signal and Information Processing, Beijing, China, 2013, pp. 422–426.
18. Dong, J.; Wang, W.; Tan, T. Casia image tampering detection evaluation database. In Proceedings of the 2013 IEEE China Summit & International Conference on Signal and Information Processing (ChinaSIP), Beijing, China, 6–10 July 2013; pp. 422–426.
19. De Carvalho, T.J.; Riess, C.; Angelopoulou, E.; Pedrini, H.; Rocha, A.D. Exposing digital image forgeries by illumination color classification. *IEEE Trans. Inf. Forensics Secur.* **2013**, *8*, 1182–1194.
20. Li, Y.; Lyu, S. Exposing DeepFake Videos by Detecting Face Warping Artifacts. *arXiv* **2019**, arXiv:1811.00656.
21. CVPRW2019_Face_Artifacts. Available online: https://github.com/danmohaha/CVPRW2019_Face_Artifacts (accessed on 16 June 2019).

