

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Face Restoration for Morphed Images Retouching

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Di Domenico, N., Borghi, G., Franco, A., Maltoni, D. (2024). Face Restoration for Morphed Images Retouching [10.1109/IWBF62628.2024.10593948].

Availability:

This version is available at: <https://hdl.handle.net/11585/967014> since: 2024-04-04

Published:

DOI: <http://doi.org/10.1109/IWBF62628.2024.10593948>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

Face Restoration for Morphed Images Retouching

Nicolò Di Domenico, Guido Borghi, Annalisa Franco, Davide Maltoni
Dipartimento di Informatica - Scienza e Ingegneria (DISI)
Università di Bologna, Cesena, Italy
{nicolo.didomenico, guido.borghi, annalisa.franco, davide.maltoni}@unibo.it

Abstract—Creating a high-quality morphed image is a laborious and time-intensive endeavor due to the necessity of manual post-processing to eliminate typical artifacts produced by landmark-based morphing techniques. At the same time, morphed images of superior quality, without noticeable visual artifacts like ghosts, noise, or blurring, present heightened success probabilities to deceive human evaluators and commercial face verification systems. Therefore, in this paper, we investigate the use of Face Restoration to automatically retouch morphed images. Specifically, we investigate the efficacy of *CodeFormer* in removing artifacts and preserving the identity of the contributing subjects. An effective retouching method would allow the generation of large datasets containing high-quality retouched morphs, even starting from existing data, that are crucial for developing and evaluating the robustness of Morphing Attack Detection (MAD) algorithms.

Index Terms—Face Morphing, Automated Artifact Retouching, Morphing Attack Detection

I. INTRODUCTION

In general, the probability of success of the Face Morphing attack [1] hinges on two key elements: the morphed image must successfully match both contributing subjects, and the morphed image should exhibit high visual quality, avoiding any visible or not-visible artifacts generated during the morphing process [2]. Indeed, these artifacts could be readily detected by a human observer, such as a police officer, or by Facial Recognition Systems (FRSs).

However, morphing algorithms can have different outcomes in terms of image quality. The majority of existing algorithms are landmark-based, employing a combination of image warping based on facial landmarks, followed by texture blending: their advantage is the simplicity in controlling the degree of similarity with the two contributing subjects, usually varying the percentage (morphing factor) of the presence of each subject in the morphing. However, landmark-based algorithms are often limited in terms of visual quality, since they may generate artifacts due to inaccuracies or insufficient precision in detecting landmark positions. In many cases, these artifacts are visible in the proximity of facial features such as the eyes, nose, and mouth [3]. Thus, manual post-processing is often

applied to morphed images, aiming to reduce or even remove any visible artifact and then to produce high-quality morphed images. However, this procedure is subjective, tedious, and time-consuming: these issues prevent the creation of large datasets with high-quality retouched morphed images that can be used to train or test Morphing Attack Detection (MAD) algorithms [4], *i.e.* automated tools explicitly devised to detect the presence of morphing in input images.

More recently, morphing algorithms based on GANs [5] have been introduced as alternatives to traditional landmark-based methods [6]–[8]. The output of these GAN-based approaches typically does not present artifacts originated from issues with the facial landmark localization but still shows artifacts specific to the GAN-based generation procedure [9]. Nonetheless, some approaches have limitations in generating high-resolution images, due to complexities in terms of memory requirements and training stability; moreover, controlling the final identity of the morphed face is more complex.

Therefore, in this paper, we investigate the use of a Face Restoration [10] approach to automatically retouch morphed images created through a landmark-based morphing approach. Specifically, we verify the capability of *CodeFormer* [11] to retouch morphed images improving the visual quality while preserving identity. A fully automated retouch operation has the potential to eliminate the necessity for manual human intervention in handling morphed images and can simplify the process of generating extensive datasets comprising high-quality images, which proves to be particularly beneficial for training and testing MAD algorithms. In addition, we believe this investigation can highlight the risk introduced by face restoration-based methods in creating morphed images that can easily deceive FRSs and humans.

II. RELATED WORK

A. Face Morphing

Face Morphing is an image manipulation technique where one image undergoes a gradual transformation into another. This technique, applied in the scenario of manipulating electronic machine-readable travel documents (eMRTD) [1], allows the creation of facial images with a dual identity. Literature studies [2], [4] report that morphed images have the capability to circumvent both FRSs and human control, and then face morphing represents a severe security threat. Moreover, the proliferation of generative AI techniques for face

This project received funding from the European Union's Horizon 2020 research and innovation program under Grant Agreement No. 883356. Disclaimer: this text reflects only the author's views, and the Commission is not liable for any use that may be made of the information contained therein.




Fig. 1: Given as input a morphed image with visible artifacts, CodeFormer [11] retouches them enhancing the visual quality.

morphing [12], [13], further amplifies this concern, streamlining the process for potential malicious actors. Additionally, the morphed images can undergo enhancement through either manual [14] or automated retouching procedures [3], effectively eradicating both discernible and imperceptible artifacts. Consequently, there exists a pressing need to develop Morphing Attack Detection (MAD) systems [4] robust to new morphing algorithms and retouching procedures.

The literature delineates two main categories of MAD methods [4], *i.e.* Single-image MAD (S-MAD) and Differential MAD (D-MAD), based on batch or incremental training [15], [16]. S-MAD systems are based exclusively on the input images embedded within electronic documents. Typically, these systems prioritize the identification of artifacts or residual traces left by the morphing process. S-MAD task is generally considered challenging, given its reliance solely on information inherent to a single input image [17]. Conversely, a D-MAD system accepts two distinct images: a live capture and a document image, which may potentially be morphed. They operate on the assumption that the live image has been procured through a trustworthy process. In this scenario, D-MAD systems not only rely on detecting morphing traces but also can exploit a comparative analysis of input identities [18].

B. Automatic Retouch of Morphed Images

Wang *et al.* [19] introduces a technique to automatically identify image manipulations, such as warping, introduced during the retouching of facial photographs through professional software. In addition, this method can reverse the warping process and restore the original image. The method can gen-

erate high-quality reconstructions, but it relies on the optical flow information of the image transformation, not available in our scenario. Seibold *et al.* presented a method [20], inspired by the style transfer approach [21], for enhancing morphed images, starting from the consideration that morphed images often exhibit fewer details compared to authentic faces. It consists of a CNN used to extract style and content features from the input images and the morphed image, respectively: then, the content is preserved to maintain the identity in the morphed image, while the resulting image’s style becomes more realistic. This approach, not based on a learning phase, does not necessitate training or fine-tuning on specific datasets. However, it is more oriented toward improving image texture rather than the explicit removal of morphing artifacts. Moreover, limitations are related to the generated image size and the inability to control the retouching areas. A method based on Conditional GAN [22] is presented in [3]: it is based on the retouching of single patches extracted from the face and then combined using a weighted blending operation. Retouched areas in images are controlled with attention maps computed through a bit-wise operation on the warped images of the two contributing subjects and fed in input to the whole system. The main limitations regard the retouching procedure based on patches, which requires the use of a non-trivial merging operation, and the impossibility of controlling the quality of the whole facial image during the generation procedure.

III. FACE RESTORATION

Face Restoration [10] is a domain-specific image restoration task, that aims to create high-quality faces from low-quality

Patch Type	CodeFormer w	Norm ↓		Difference ↓		RMSE ↓			δ -metrics ↑			Indexes ↑	
		L_1	L_2	Abs	Sqr	Lin	Log	Scl	1.25	1.25 ²	1.25 ³	PSNR	SSIM
Right Eye	0.00	10.16	6673	0.11	2.24	15.0	0.32	0.54	2.48	2.83	2.93	24.76	0.69
	0.50	9.12	5966	0.10	1.85	13.5	0.31	0.50	2.54	2.85	2.94	25.74	0.72
	1.00	8.70	5840	0.10	1.80	13.2	0.30	0.50	2.57	2.86	2.94	26.01	0.73
Left Eye	0.00	10.60	6917	0.11	2.33	15.6	0.31	0.52	2.48	2.83	2.93	24.47	0.68
	0.50	9.51	6170	0.10	1.90	13.9	0.30	0.48	2.57	2.86	2.95	25.50	0.71
	1.00	9.02	6000	0.10	1.83	13.5	0.30	0.49	2.59	2.87	2.95	25.83	0.72
Nose	0.00	10.48	6741	0.11	2.25	15.2	0.61	1.05	2.55	2.85	2.93	24.71	0.73
	0.50	9.36	6001	0.10	1.83	13.5	0.58	0.99	2.63	2.88	2.94	25.76	0.76
	1.00	8.40	5406	0.09	1.55	12.2	0.66	1.13	2.67	2.90	2.95	26.71	0.79
Mouth	0.00	9.20	5545	0.09	1.54	12.5	0.23	0.38	2.66	2.91	2.97	26.47	0.71
	0.50	8.49	5056	0.08	1.28	11.4	0.22	0.36	2.72	2.93	2.97	27.23	0.74
	1.00	7.88	4750	0.08	1.17	10.7	0.21	0.35	2.74	2.94	2.98	27.84	0.75
All Patches	0.00	10.03	6389	0.11	2.04	14.4	0.35	0.59	2.55	2.86	2.94	25.22	0.70
	0.25	9.36	5953	0.10	1.79	13.4	0.34	0.56	2.61	2.88	2.95	25.83	0.72
	0.50	9.07	5734	0.10	1.68	12.9	0.33	0.55	2.62	2.89	2.95	26.16	0.74
	0.75	8.79	5573	0.09	1.60	12.6	0.34	0.56	2.64	2.89	2.96	26.43	0.74
	1.00	8.45	5439	0.09	1.55	12.3	0.35	0.58	2.65	2.89	2.96	26.70	0.75

TABLE I: Pixel-wise metrics computed on retouched images from FRGC_S [3] dataset through CodeFormer [11].

input images. Depending on the elements that are responsible for the low quality of input, a variety of subtasks is addressed in the literature: face deblurring [23] (sharp images from blurry ones), face denoising [24] (removing the noise), face super-resolution [25] (enhancing the spatial resolution of input facial image) and face artifact removal [26] (removal of artifacts introduced due to a lossy compression). Whether the restoration is not based on the knowledge about the type of degradation present in the input image, the task is referred to as *blind* restoration [11]. At the time of writing, no studies focus on the use of face restoration for morphing artifact removal.

CodeFormer [11] is a Transformer-based prediction network that models the global composition of natural faces. This blind face restoration approach relies on the concept of vector quantization and the whole system is divided in different components. A quantized autoencoder is used to obtain a discrete codebook, from which a related decoder generates the self-reconstruction of the input image. The subsequent face restoration procedure is based on the prior from the codebook combination and decoder. Then, a transformer is exploited to predict a code combination starting from a low-quality input. It is worth noting that the restoration procedure is controlled by a parameter $w \in [0, 1]$ that controls the trade-off between *quality* and *fidelity* of the final image. Generally, a small w reduces the reliance of the model on the input image, allowing significant changes in the generated one. Conversely, larger w preserves the information from the input image, reducing the impact of changes but improving the fidelity. In our experiments, we use the weights of the official implementation¹; the face detection is performed through RetinaFace [27], and the model has a codebook size of 1024, embedding size of 512, and the number of heads and layers is 8 and 9, respectively.

¹<https://github.com/sczhou/CodeFormer>

IV. EXPERIMENTAL EVALUATION

A. Datasets

FRGC_S [3]. This dataset simulates morphed images before and after a retouching procedure: indeed, each morphed image is obtained using two faces belonging to the same subject taken from the FRGC dataset [28]: therefore, the difference between the contributing and the morphed images is related only to artifacts and not to different identities. The presence of artifacts is further improved by applying a perturbation on the detected landmarks of the first image used for the morphing. This dataset is useful for the quantitative evaluation of the retouching capability, thanks to the presence of the reference image on which it is possible to compute pixel-based metrics.

FRGC_M [3]. Differently from FRGC_S, this dataset consists of real morphed images, *i.e.* images obtained combining two different identities. The dataset consists of 1060 morphed images and 32 subjects. Each subject is combined with all other subjects of the same gender. This results in a dataset containing morphed images with varying quality levels. FRGC_M dataset is useful to carry out experiments related to identity preservation and MAD testing.

SMDD [29]. The Synthetic Morphing Attack Detection Development dataset contains synthetic facial images exploiting the StyleGAN2-ADA model [30] trained on Flickr-Faces-HQ Dataset (FFHQ) dataset [31]. Synthetic faces are then used to create morphed images through a widely used morphing algorithm [32]. We exploit the SMDD dataset to assess the retouching ability of CodeFormer varying the source data and the algorithm used for morphing with respect to FRGC_M.

B. Evaluation of visual quality

To evaluate the visual quality, we exploit several pixel-wise metrics to check the distance of generated images with respect to the ground truth reference. Following [33], [34], metrics are the L_1 and L_2 distances, the Root Mean Square Error

Patch Type	Method	Norm ↓		Difference ↓		RMSE ↓			δ -metrics ↑			Indexes ↑	
		L_1	L_2	Abs	Sqr	Lin	Log	Scl	1.25	1.25 ²	1.25 ³	PSNR	SSIM
Right Eye	CodeFormer [3]	10.38	6664	0.11	2.37	15.0	0.30	0.42	2.49	2.83	2.93	25.84	0.77
		8.70	5840	0.10	1.80	13.2	0.30	0.50	2.57	2.86	2.94	26.01	0.73
Left Eye	CodeFormer [3]	10.35	6627	0.11	2.23	14.9	0.32	0.51	2.51	2.84	2.94	25.47	0.74
		9.02	6000	0.10	1.83	13.5	0.30	0.49	2.59	2.87	2.95	25.83	0.72
Nose	CodeFormer [3]	8.00	4939	0.08	1.09	11.1	0.42	0.71	2.79	2.95	2.98	27.72	0.85
		8.40	5406	0.09	1.55	12.2	0.66	1.13	2.67	2.90	2.95	26.71	0.79
Mouth	CodeFormer [3]	8.48	5344	0.08	1.24	12.1	0.20	0.30	2.76	2.95	2.98	26.98	0.81
		7.88	4750	0.08	1.17	10.7	0.21	0.35	2.74	2.94	2.98	27.84	0.75
All Patches	CodeFormer [3]	9.25	5861	0.09	1.70	13.2	0.30	0.47	2.65	2.90	2.96	26.52	0.79
		8.45	5439	0.09	1.55	12.3	0.35	0.58	2.65	2.89	2.96	26.70	0.75

TABLE II: Pixel-wise metrics computed on retouched images using FRGC_S dataset [3].

(RMSE), and three δ -metrics, representing the percentage of pixels falling below a specified threshold. Furthermore, our analysis incorporates the Peak Signal-to-Noise Ratio (PSNR), contributing to a comprehensive evaluation of image quality. We also use the Structural Similarity (SSIM), which estimates the perceived visual similarity of two images.

To maintain the consistency with previous literature evaluation, all pixel-wise metrics are computed on color patches with a spatial resolution of 256×256 . Moreover, the use of patch-based evaluation allows to focus the attention on facial areas generally affected by artifacts, such as the eyes (pupils and eye corners), nose and mouth. Being aware the assessment of the visual quality of generated images is still a challenging task [35], methods are tested on a list² of images belonging to FRGC_S that present the most visible artifacts, to highlight the retouching ability of the analyzed methods.

The first step is the investigation of the impact of different weight values (w), that control the trade-off between quality and fidelity of CodeFormer. This investigation requires the use of the FRGC_S dataset, and pixel-wise metrics to quantify the difference between the ground truth and the restored morphed image. Results are reported in Table I, in which the first lines show the metrics for each patch type, and the last line reports the averaged values computed on all patches. We observe that the use of high values allows an effective face restoration, despite the presence of strong visible artifacts in FRGC_S images. All metrics concur in highlighting that low values of w yield excessively significant variations in the output images.

As the second step, we compare CodeFormer with a competing retouching method of morphed images available in the literature [3]. The results, also in this case reported for the different patch types and then finally averaged, are shown in Table II. For CodeFormer, we report the best combination ($w = 1$), as determined in the previous analysis. From a general point of view, the majority of the metrics reveal the superior efficacy of CodeFormer in image retouching, especially for eye patches. In particular, the PSNR, which discerns the presence of noise in images and is measured on a logarithmic scale, exhibits superior performance.

²<https://miatbiolab.csr.unibo.it/public-resources>

Finally, in Figure 1 some output samples are reported, in which it is possible to observe that CodeFormer effectively handles accessories, such as glasses, but also the non-trivial trade-off between the ability to correct even strong artifacts (such as the mouth in the third line) and the faculty to preserve the same identity, as discussed in the following.

C. Evaluation of Identity Preservation

For the evaluation of identity preservation, we employ the Mated Morph Presentation Match Rate (MMPMR) metric, introduced in [36] as a gauge of the vulnerability of FRSS. Indeed, MMPMR denotes the ratio of morphed images that can be matched with both contributing subjects. In our experiments, we employ a commercial SDK, VeriLook v12 by Neurotechnology³ as face verification tool. Then, each morphed image is compared against a test image (distinct from the one used for morphing creation) of each contributing subject.

³<https://www.neurotechnology.com/verilook.html>




Fig. 2: MMPMR measured in the identity preservation test for different w of CodeFormer and similarity scores of the contributing subjects (see Sect. IV-C). GAN refers to [3].




Fig. 3: (a) DET curve of the R-3 model [17] tested on SMDD and $FRGC_M$ datasets, before and after the retouching procedure. (b) DET curve of the R-3 model fine-tuned on the retouched images and tested on SMDD and $FRGC_M$.

Adhering to the Frontex guidelines for face verification at ABC gates, the threshold for both SDKs was set to operate at a False Acceptance Rate (FAR) of 0.1%.

This evaluation is conducted on the real morphed images of the $FRGC_M$ dataset for both our model and [3], and the results are reported in Figure 2 as a function of the similarity score between the contributing subject and different values of w . We observe that for a wide range of similarity scores ($[10, 40]$) the automated retouching with the optimal $w = 1$ has no impact on the MMPMR, thereby confirming the preservation of the morphed identity. Only for very low similarity scores (< 10), a limited decrease of about 4% in MMPMR is observed. Indeed, as reported in [3], for lower similarity scores the likelihood of success is considerably restricted, and even minor manual or automated modifications to the image frequently result in falling below the predetermined matching threshold. As expected, lower values of w significantly affect the resulting identity, with a drop of performance of about 20% with high similarity scores ($[30, 40]$), and more than 60% with low similarity ones ($[0, 10]$).

D. Impact on the S-MAD task

In this section, we evaluate the impact of the retouching procedure both on S-MAD training and testing through the error-based metrics tailored to the MAD task [4]. The Bona Fide Presentation Classification Error Rate (BPCER) quantifies the rate of genuine images incorrectly identified as morphed, whereas the Attack Presentation Classification Error Rate (APCER) measures the rate of morphed images inaccurately classified as bona fide. BPCER values are presented alongside predetermined APCER values, *i.e.* $BPCER_{0.1}$, $BPCER_{0.05}$, and $BPCER_{0.01}$, which represent the minimum achievable BPCER with an APCER not exceeding 10%, 5%, and 1%, respectively. All these metrics and values are condensed in the Detection Error Trade-off (DET) curve.

For these experiments, we retouch with CodeFormer ($w = 1$) the morphed images of $FRGC_S$, $FRGC_M$, and a subset of

1000 images of SMDD, obtaining $FRGC_S^{CF}$, $FRGC_M^{CF}$, and $SMDD^{CF}$, respectively.

Firstly, we examine whether retouched images possess the capability to deceive an S-MAD algorithm. For the S-MAD model, aware that several models are available in the literature, we implement the current state-of-the-art solution on the FVC-onGoing platform (SOTAMD_D benchmark), referred to as “R-3”⁴ in [17]. This MAD solution is based on the Inception-ResNet architecture [37], and input faces are cropped through the MTCNN [38] face detector and belong to a variety of training datasets, including both landmark- and GAN-based morphing algorithms [17]. The performance of R-3 is in Figure 3a, tested on both $FRGC_M$ (blue line), $FRGC_M^{CF}$ (blue dotted line), SMDD (orange line), and $SMDD^{CF}$ (orange dotted line): a significant increase in the error rates denotes a superior quality of retouched images, and a reduced accuracy of the S-MAD model that misclassifies morphed as bona fide.

Secondly, we fine-tune (5 epochs with a learning rate of 10^{-4} and SGD optimizer with 0.9 of momentum) the same R-3 method on $FRGC_S^{CF}$, obtaining the blue curves depicted in Figure 3b; we repeat the process fine-tuning on $SMDD^{CF}$ for 1 epoch to obtain the orange curves; results reveal that including the retouched images in the training dataset represent a suitable and effective solution to significantly reduce the error rates on retouched morphed images, at the cost of a limited error increment on no-retouched ones.

Considering both experiments, we observe that the impact of retouched images on the S-MAD performance is significant: indeed, retouched images exhibit a higher attack potential representing a viable tool to evade detection by S-MAD models. In addition, we suggest that it is important to include them, possibly created through a variety of face restoration methods, in training data to improve the robustness of the deployed model against new approaches able to enhance the quality of morphing attacks.

V. CONCLUSION

In this paper, we have introduced the use of the blind face restoration task, and specifically the CodeFormer method, to automatically remove artifacts produced by landmark-based morphing algorithms. Quantitative and qualitative results are promising, even though the use of w is thorny: indeed, lower values of w remove even strong artifacts, but tend to modify the final identity, and vice-versa. Plenty of future works is planned, including the application of the retouching procedure through new face restoration models on different initial data, on images produced with various morphing algorithms. Moreover, considering the importance of including retouched images in face-morphing training datasets, we plan to release a retouched version of public morphed datasets. In addition, an analysis of new face restoration models will be conducted, to monitor the rise of new security threats and to develop proper countermeasures. Finally, it is interesting to include also an ISO/ICAO compliance check to analyze if generated images are still suitable for electronic identity documents.

⁴<https://github.com/ndido98/ubo-smad-r3>

ACKNOWLEDGMENTS

We acknowledge the CINECA award under the IS CRA initiative, for the availability of high-performance computing resources and support. We thank Andrea Pilzer from the NVIDIA AI Technology Center for his support.

REFERENCES

- [1] M. Ferrara, A. Franco, and D. Maltoni, "The magic passport," in *IEEE International Joint Conference on Biometrics (IJCB'14)*, 2014, pp. 1–7.
- [2] U. Scherhag, A. Nautsch, C. Rathgeb, M. Gomez-Barrero, R. N. Veldhuis, L. Spreeuwens, M. Schils, D. Maltoni, P. Grother, S. Marcel *et al.*, "Biometric systems under morphing attacks: Assessment of morphing techniques and vulnerability reporting," in *2017 International Conference of the Biometrics Special Interest Group (BIOSIG)*, 2017.
- [3] G. Borghi, A. Franco, G. Graffieti, and D. Maltoni, "Automated artifact retouching in morphed images with attention maps," *IEEE Access*, 2021.
- [4] K. Raja, M. Ferrara, A. Franco, L. Spreeuwens, I. Batskos, F. de Wit, M. Gomez-Barrero, U. Scherhag, D. Fischer, S. K. Venkatesh *et al.*, "Morphing attack detection-database, evaluation platform, and benchmarking," *IEEE transactions on information forensics and security*, vol. 16, pp. 4336–4351, 2020.
- [5] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proceedings of the 27th International Conference on Neural Information Processing Systems-Volume 2*, 2014, pp. 2672–2680.
- [6] N. Damer, A. M. Saladié, A. Braun, and A. Kuijper, "Morgan: Recognition vulnerability and attack detectability of face morphing attacks created by generative adversarial network," in *2018 IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, 2018, pp. 1–10.
- [7] H. Zhang, S. Venkatesh, R. Ramachandra, K. Raja, N. Damer, and C. Busch, "Mipgan – generating robust and high quality morph attacks using identity prior driven gan," 2020.
- [8] N. Damer, K. Raja, M. Süßmilch, S. Venkatesh, F. Boutros, M. Fang, F. Kirchbuchner, R. Ramachandra, and A. Kuijper, "Regenmorph: Visibly realistic gan generated face morphing attacks by attack regeneration," in *Advances in Visual Computing: 16th International Symposium, ISVC 2021, Virtual Event, October 4-6, 2021, Proceedings, Part I*. Springer, 2021, pp. 251–264.
- [9] L. Chai, D. Bau, S.-N. Lim, and P. Isola, "What makes fake images detectable? understanding properties that generalize," in *European Conference on Computer Vision*. Springer, 2020, pp. 103–120.
- [10] T. Wang, K. Zhang, X. Chen, W. Luo, J. Deng, T. Lu, X. Cao, W. Liu, H. Li, and S. Zafeiriou, "A survey of deep face restoration: Denoise, super-resolution, deblur, artifact removal," *arXiv preprint arXiv:2211.02831*, 2022.
- [11] S. Zhou, K. Chan, C. Li, and C. C. Loy, "Towards robust blind face restoration with codebook lookup transformer," *Advances in Neural Information Processing Systems*, vol. 35, pp. 30 599–30 611, 2022.
- [12] H. Zhang, S. Venkatesh, R. Ramachandra, K. Raja, N. Damer, and C. Busch, "Mipgan—generating strong and high quality morphing attacks using identity prior driven gan," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 3, no. 3, pp. 365–383, 2021.
- [13] S. Venkatesh, H. Zhang, R. Ramachandra, K. Raja, N. Damer, and C. Busch, "Can gan generated morphs threaten face recognition systems equally as landmark based morphs?-vulnerability and detection," in *2020 8th International Workshop on Biometrics and Forensics (IWBF)*. IEEE, 2020, pp. 1–6.
- [14] M. Ferrara, A. Franco, and D. Maltoni, "Face demorphing," *IEEE Transactions on Information Forensics and Security*, 2017.
- [15] L. Pellegrini, G. Borghi, A. Franco, and D. Maltoni, "Detecting morphing attacks via continual incremental training," in *2023 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2023.
- [16] G. Borghi, G. Graffieti, A. Franco, and D. Maltoni, "Incremental training of face morphing detectors," in *2022 26th International Conference on Pattern Recognition (ICPR)*. IEEE, 2022, pp. 914–921.
- [17] G. Borghi, N. Di Domenico, A. Franco, M. Ferrara, and D. Maltoni, "Revelio: A modular and effective framework for reproducible training and evaluation of morphing attack detectors," *IEEE Access*, 2023.
- [18] N. Di Domenico, G. Borghi, A. Franco, and D. Maltoni, "Combining identity features and artifact analysis for differential morphing attack detection," in *International Conference on Image Analysis and Processing*. Springer, 2023, pp. 100–111.
- [19] S.-Y. Wang, O. Wang, A. Owens, R. Zhang, and A. A. Efros, "Detecting photoshopped faces by scripting photoshop," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- [20] C. Seibold, A. Hilsman, and P. Eisert, "Style your face morph and improve your face morphing attack detector," in *2019 International Conference of the Biometrics Special Interest Group (BIOSIG)*. IEEE, 2019, pp. 1–6.
- [21] L. Gatys, A. Ecker, and M. Bethge, "A neural algorithm of artistic style," *Journal of Vision*, vol. 16, no. 12, pp. 326–326, 2016.
- [22] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017.
- [23] R. Yasarla, F. Perazzi, and V. M. Patel, "Deblurring face images using uncertainty guided multi-stream semantic networks," *IEEE Transactions on Image Processing*, vol. 29, pp. 6251–6263, 2020.
- [24] E. Luo, S. H. Chan, and T. Q. Nguyen, "Adaptive image denoising by targeted databases," *IEEE transactions on image processing*, 2015.
- [25] E. Zhou, H. Fan, Z. Cao, Y. Jiang, and Q. Yin, "Learning face hallucination in the wild," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 29, no. 1, 2015.
- [26] L. Yang, S. Wang, S. Ma, W. Gao, C. Liu, P. Wang, and P. Ren, "Hifacegan: Face renovation via collaborative suppression and replenishment," in *Proceedings of the 28th ACM international conference on multimedia*, 2020, pp. 1551–1560.
- [27] S. I. Serengil and A. Ozpinar, "Hyperextended lightface: A facial attribute analysis framework," in *2021 International Conference on Engineering and Emerging Technologies*. IEEE, 2021, pp. 1–4.
- [28] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek, "Overview of the face recognition grand challenge," in *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, vol. 1. IEEE, 2005, pp. 947–954.
- [29] N. Damer, C. A. F. López, M. Fang, N. Spiller, M. V. Pham, and F. Boutros, "Privacy-friendly synthetic data for the development of face morphing attack detectors," pp. 1605–1616, 2022.
- [30] T. Karras, M. Aittala, J. Hellsten, S. Laine, J. Lehtinen, and T. Aila, "Training generative adversarial networks with limited data," in *Proc. NeurIPS*, 2020.
- [31] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4401–4410.
- [32] M. Satya. Face morph using opencv — c++ / python. [Online]. Available: <https://learnopencv.com/face-morph-using-opencv-cpp-python/>
- [33] D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," *Advances in Neural Information Processing Systems*, vol. 27, pp. 2366–2374, 2014.
- [34] S. Pini, G. Borghi, and R. Vezzani, "Learn to see by events: Color frame synthesis from event and rgb cameras," in *International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, vol. 4, 2020, pp. 37–47.
- [35] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved techniques for training gans," *arXiv preprint arXiv:1606.03498*, 2016.
- [36] U. Scherhag, A. Nautsch, C. Rathgeb, M. Gomez-Barrero, R. N. J. Veldhuis, L. Spreeuwens, M. Schils, D. Maltoni, P. Grother, S. Marcel, R. Breithaupt, R. Ramachandra, and C. Busch, "Biometric systems under morphing attacks: Assessment of morphing techniques and vulnerability reporting," in *2017 International Conference of the Biometrics Special Interest Group (BIOSIG)*, 2017, pp. 1–7.
- [37] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9.
- [38] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE signal processing letters*, vol. 23, no. 10, pp. 1499–1503, 2016.