MONOT: High-Quality Privacy-compliant Morphed Synthetic Images for Everyone

Guido Borghi[1,2], Nicolò Di Domenico[1], Matteo Ferrara[1], Annalisa Franco[1], Uzma Latif[1], Davide Maltoni[1]
[1]Department of Computer Science and Engineering, University of Bologna, Italy
[2]Department of Education and Humanities, University of Modena and Reggio Emilia, Italy
guido.borghi@unimore.it, {name.surname}@unibo.it, uzma.latif2@unibo.it

***Abstract*—Morphing Attack Detection (MAD) is a critical task in biometric security, aimed at identifying and mitigating the risks posed by morphing attacks, where a face image is manipulated to resemble multiple individuals. Therefore, MAD systems are essential to prevent unauthorized access and ensure the integrity of biometric authentication procedures. However, the acquisition, storing and transfer of real biometric data on which they are based are limited by ethical, legal, and privacy concerns, and this hinders their accuracy. To address these issues, we propose and release MONOT, a new dataset of synthetic morphed images. The dataset includes high-quality synthetic morphed images that are ISO/ICAO compliant and have the characteristics of real biometric data without compromising individual privacy. The morphing procedure is applied through six different morphing algorithms, providing a great level of data variability. Our experimental results demonstrate MONOT morphed images show a high attack potential and that MAD systems trained on MONOT exhibit high detection performance across various morphing techniques. All these elements highlight the dataset's effectiveness in supporting the development of robust and generalized MAD systems.**

***Index Terms*—Morphing Attack, Morphing Attack Detection, Synthetic Data, Face-based Verification Methods**

I. Introduction

The accuracy of biometric systems is critical for ensuring secure and reliable authentication in various applications, including access control, identity verification, and border security. Biometric systems leverage unique physiological or behavioral characteristics, such as fingerprints, iris patterns, and facial features, to identify individuals. The precision of these systems directly impacts their ability to distinguish between legitimate users and imposters, thereby preventing unauthorized access and enhancing security.

However, despite their high accuracy and widespread adoption, biometric systems are vulnerable to recently discovered attacks, one of which is the so-called Morphing Attack [1]. Specifically, the morphing attack involves the creation of a single image that combines the facial features of two individuals, enabling both to be authenticated as the same person. This type of attack poses significant security risks, especially in the context of travel and identity documents, where it can facilitate illegal activities such as unauthorized border crossings through the Automatic Border Control (ABC) gates in international airports and identity fraud [2].

(a) (b) (c)

Fig. 1: Example of the Face Morphing attack: two individuals (1a and 1c) are combined to produce a morphed image (1b). This hybrid image conceals the identity of a criminal within the image of an accomplice, and it can deceive both human examiners and automated face verification systems [2].

Unfortunately, detecting morphing attacks is a challenging task due to the subtlety of the image modifications, necessitating advanced detection mechanisms. To address this challenge, the development of robust Morphing Attack Detection (MAD) systems [3] is essential and strongly needed.

However, the creation and sharing of real biometric data for training and evaluating these systems are fraught with privacy and legal issues. Indeed, the use of personal data, especially if consists of faces, is limited by privacy concerns, consent requirements, and stringent regulations, which mandate the protection of personal data [4]. As a consequence, these constraints limit the availability of diverse and representative datasets needed for effective MAD development.

In this context, synthetic data emerges as a viable alternative [5] since new datasets can be generated to mimic the characteristics of real biometric data without compromising individual privacy. Synthetic data has garnered significant interest within the scientific community, as recent methods (*e.g.* Diffusion Models [6] and Variational Autoencoders [7]) for generating highly realistic data have been introduced. In other words, synthetic data offers several advantages, including the ability to generate large, diverse datasets that can enhance the training and testing of MAD systems.

Therefore, in this paper, we propose a new synthetic dataset, namely MONOT[1], specifically designed for the morphing

This project received funding from the European Union's Horizon 2020 research and innovation program under Grant Agreement No. 883356. 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.

[1]https://miatbiolab.csr.unibo.it/monot-synthetic-dataset

attack detection task. MONOT dataset consists of a collection of morphed images, created through a variety of morphing algorithms. Morphed images exhibit high visual quality in terms of ISO/ICAO compliance [8], reproducing the operational scenario in which they are used as document photos. In addition, for each subject, several "in the wild images" are provided resembling, for instance, the live acquisition at the airport gates. Concluding, the MONOT dataset aims to provide a comprehensive and privacy-friendly resource for training and evaluating MAD systems.

II. Related Work

A. Synthetic Data for Morphing Attack Detection

The emergence of advanced AI-based generative algorithms has led to the creation of numerous synthetic datasets in the literature. In this context, a significant number of these datasets [9]–[13] are specifically designed for face recognition tasks. Consequently, they often focus on including a high diversity of identities and head poses, while potentially overlooking certain standard requirements [14]. Other datasets utilize recent diffusion models [15], [16], which aim to maintain and diversify identities by inverting pre-trained face recognition models, resulting in face images with greater realism compared to GANs [17], [18]. A smaller subset of synthetic datasets is developed for alternative tasks, such as age and gender recognition [19] or the morphing attack detection [20] task investigated in this paper.

Specifically, for the MAD task, one of the most important synthetic datasets is the Synthetic Morphing Attack Detection Development dataset (SMDD) [20]. It consists of about 30k attack and 50k bona fide images, and it is a valuable resource for training and evaluating morphing attack detection systems. By providing a high volume of diverse synthetic data, the SMDD dataset enables more robust and generalized training of MAD systems, which can lead to improved detection performance even on previously unseen attack types. In relation to our work, we observe the visual quality of the images in SMDD is variable. In particular, all the morphed images present a not uniform background, since the starting face images are generated "in the wild". This condition is not very realistic, since morphed images are used as document photos, and then they must be compliant with the ISO/ICAO guidelines [8], [21]. These guidelines, developed both by ISO and ICAO institutions, play a crucial role in numerous biometric applications. Indeed, meeting these standards for facial images in official documents greatly improves face verification accuracy [22], [23]. Compliant images ensure consistency in quality, thereby facilitating more precise matching. This uniformity is essential in reducing the risks of false positives and false negatives, which are critical for maintaining the reliability of biometric systems. Among the requirements, ISO/ICAO guidelines impose that document images have a uniform and light background, a frontal head and shoulder pose of the acquired subject, and a neutral expression, all elements not always available in the SMDD dataset as depicted in Figure 2. Starting from these observations, MONOT consists of ISO/ICAO compliant morphed images (see Fig. 3), and then MAD systems can be trained and evaluated with more realistic synthetic images.

Fig. 2: Samples of morphed images belonging to the SMDD dataset [20]. In these cases, the realism of these images is limited, since morphed images used in documents should be compliant with ISO/ICAO guidelines [8].

B. The Morphing Attack

As mentioned, face morphing is an image manipulation technique used to progressively transform one face into another. Initially described in [1] in the context of electronic machine-readable travel (eMRTD) documents, this technique allows for the creation of hybrid faces with dual identities, capable of evading both automated face verification systems and human examiners [2], [3]. Recently, the application of this technique has expanded to include 3D data [24].

The attack is further complicated by the rise of generative AI techniques for face morphing [25]–[27], which have made the process more accessible to malicious actors. These morphed images can be further enhanced through manual [28] or automated retouching [29], [30], effectively removing both visible and subtle artifacts, thus increasing the difficulty of detection. Consequently, there is an urgent need to develop new MAD systems. These automated tools must be explicitly designed to detect morphing in facial images with high accuracy and the ability to generalize to previously unseen images.

C. Morphing Attack Detection Models

The existing literature primarily categorizes MAD methods into two types [3]: Single-image MAD (S-MAD) and Differential MAD (D-MAD).

S-MAD systems [31]–[33] focus on artifacts or traces left by morphing procedures in input images. Typically, these systems rely solely on the information provided by the single input images [34]. This task is generally considered challenging because it typically relies only on information available in a single image, as evidenced by experimental results [35].

On the other hand, D-MAD systems [36]–[39] take two different images as input: one from a trusted live capture and another from the document: this image could potentially be morphed. These systems operate on the assumption that at

Fig. 3: Samples of morphed images of the MONOT dataset. As shown, all images present a high visual quality and are ISO/ICAO compliant [8]. Among the others, all images have a uniform and light background, and the subject depicted has a neutral expression and uniform face illumination.

least one of these inputs has been acquired through a reliable process, such as from a camera at an ABC gate or through a procedure supervised by a law enforcement officer.

Focusing on S-MAD, early efforts in MAD focused on handcrafted features and classical machine learning algorithms. For instance, some studies employed texture descriptors [40], [41] to differentiate between morphed and bona fide images. Usually, these methods often struggled with generalization to new attack types and variations in image quality. More recent advancements have leveraged deep learning techniques [33], [42], which have shown superior performance in detecting morphing attacks.

III. MONOT Dataset

MONOT dataset consists of about 15k synthetic morphed images specifically generated for the morphing attack detection task using six different state-of-the-art morphing approaches and two morphing factors (see Sect. III-A).

For the generation of the dataset, we utilized synthetic faces from the ONOT dataset [43], a synthetic collection of high-quality face images designed to adhere to the ISO/IEC 39794-5 standards [8] and the guidelines of the International Civil Aviation Organization (ICAO) for electronic Machine-Readable Travel Documents (eMRTD).

The presence of ISO/ICAO compliant images is important since, in this work, we made a significant effort to reproduce as accurately as possible the real scenario. Indeed, in the morphing attack, the morphing procedure is applied to document images. As mentioned, these facial images, in order to be included in official documents, must follow strict quality requirements defined by ISO and ICAO institutions. Therefore, the development of datasets containing morphed images for the eMRTD use case should produce strictly ISO/ICAO compliant images, to increase the level of realism and efficacy of the proposed solutions. This is an added value of this work because, to the best of our knowledge, most of the existing synthetic face datasets have not been explicitly developed for this application and are not realistic candidates to simulate eMRTD images.

For each of the 254 subjects of the ONOT dataset, we provide also ten other synthetic images taken "in the wild", *i.e.* without adhering to the ISO/ICAO guidelines. In this way, we aim to simulate a real-world operational scenario at the airport gates, in which the image stored into the document (bona fide or morphed) is compared with a live acquisition, that indeed is collected in a low-controlled scenario (*e.g.*, the background, as well the illumination, can be not uniform). These additional images, referred to as "gate images", are created through the approach described in Section III-B.

A. Morphed Images Generation

The morphed images have been created by selecting the morph pairing candidates with high comparison scores from three Commercial-Off-The-Shelf (COTS) Face Recognition Systems (FRSs). Similar to [28], the selection of candidate images to produce morphing cases is performed as follows. An image of each subject (*i.e.*, the criminal) is compared with the image of other subjects of the same gender and ethnicity (*i.e.*, possible accomplices). Since three different FRSs have been involved, a unique value $v_{i,j}$ to select the most promising accomplices is computed as:

$$v_{i,j} = \frac{1}{3} \cdot \sum_{k=1}^{3} \frac{\tau_k - s_{i,j}^k}{\tau_k}$$

where $s_{i,j}^k$ is the similarity score between subjects i and j provided by FRS k and τ_k is the score threshold suggested by the k^{th} FRS corresponding to a False Acceptance Rate (FAR) equals to 0.1%. Value $v_{i,j}$ represents how far from the FAR=0.1% thresholds the verification scores are on the average; lower values indicate more similar subjects.

Given the criminal subject i, the candidate accomplices j are sorted in increasing order of $v_{i,j}$ and the top five candidates are selected as accomplices. This choice is aimed at maximizing the probability of fooling FRSs at the gate. Note that, none of the selected pairs were able to fool the three FRSs at the same time. Since each of the 254 subjects is paired with five other subjects, a total of 1270 morphing pairs have been selected.

To increase the diversity and the challenging nature of the dataset in order to simulate a realistic scenario, each image pair has been morphed using six different state-of-the-art landmark-based morphing algorithms (*i.e.*, C01 [44], C02 [45], C03 [3], C05 [28], [46], C015 [47], [48], and C016 [49]) and two morphing factors of 0.3 and 0.5 obtaining a total of $1270 \times 6 \times 2 = 15240$ morphed images. Figure 3 and 4 shows an example of morphed images generated using the six morphing algorithms used to create the MONOT dataset.

B. Gate Images Generation

For the generation of synthetic gate images, we focus on the Arc2Face model [50], an identity-conditioned face foundation model that can generate diverse photo realistic images from the ArcFace [51] embedding of an individual. As reported in [50], this model achieves a higher level of facial similarity compared to existing generative models.

Fig. 4: Example of morphed images contained in the MONOT dataset obtained using six morphing algorithms (see Sect. III-A) with a morphing factor of 0.5 (first row) and 0.3 (second row) on the two subjects reported in Figures 1a and 1c as criminal and accomplice, respectively.

Fig. 5: Example of the gate images generated through the Arc2Face [50] method. In the first column, the original subject image is reported, while in the remaining columns are reported the corresponding gate images.

Arc2Face is built upon a pre-trained Stable Diffusion model but is specifically adapted for the task of ID-to-face generation, conditioned solely on identity vectors. Unlike recent approaches that combine identity with text embeddings for zero-shot personalization of text-to-image models, this method emphasizes the compactness of face recognition features, which can fully capture the essence of the human face without the need for hand-crafted prompts. This allows Arc2Face to generate images using only the identity features, providing a robust prior for tasks where ID consistency is crucial.

In this context, our primary objective is to generate new realistic images of an individual to simulate images captured live in an unconstrained scenario such as an ABC gate (see Figure 5). All the generated images have been created following the official guidelines and using the default hyper-parameters.

IV. Experimental evaluation

A. Datasets

Progressive Morphing Database (PMDB) [28]: it is a collection of 1108 morphed images generated by applying the C05 morphing algorithm [28], [46] to AR [52], FRGC [53], and Color Feret [54] datasets.

Idiap Morph [55], [56]: it is a collection of datasets created using different morphing algorithms. We considered here the portion of images generated with FaceMorpher [44] (algorithm C01), which are of suboptimal visual quality due to the presence of artifacts in both the background and foreground.

FEI Morph [57]: it is a dataset generated using the images contained in the *FEI Face Database* [58], which includes 200 subjects equally distributed between male and female subjects, with an age mainly ranging from 19 to 40 years. The images are characterized by a good variability in terms of appearances, hairstyles, and presence of accessories. This dataset contains 6k morphed images obtained with three different morphing algorithms, namely C02 [45], C03 [59], and C08 [59], employing two different morphing factors (0.3 and 0.5).

ChiMo [33] : it is a compilation of morphed images generated from the images with neutral expressions of the Chicago

Faces Database (CFD) [60] which includes photographs of 831 individuals from diverse ethnies. Also for this dataset, the morphed images used in the experiments have been generated by the C02 [45], C03 [59], and C08 [59], morphing algorithms. The specific use of the datasets in the different experiments will be described in the related section.

B. Results on MONOT attack potential

A first analysis has been conducted to evaluate the attack potential of the MONOT morphed images, and to compare it to the attack potential of a real dataset of morphed images. For this evaluation, we used the Morphing Attack Potential (MAP) metric recently introduced in [61] aimed at quantifying the attack potential of a dataset M of morphed images analyzing the combined impact of a variable number of probe images and multiple FRSs. MAP is defined as a matrix with a number of rows corresponding to the number of probe gate images for each morphed image and a number of columns equal to the number of FRSs considered in the evaluation. A generic element of the matrix MAP$[r, c]$ represents the number of morphed images in the dataset M that can be successfully matched with at least r probe images by at least c FRSs. This definition implies that the MAP values naturally decrease as we move towards the bottom-right corner that refers to the most dangerous images in the dataset.

For MAP computation, three COTS FRSs (referred as FRS_1, FRS_2 and FRS_3) which provided top performance in the "Face Recognition Vendor Test (FRVT)—1:1 Verification" [62] are considered. The verification thresholds have been fixed as suggested by the three FRSs to work at a FAR=0.1%, which is the reference value for face recognition in eMRTD. Table I reports the comparison, in terms of MAP, between MONOT and FEI for different morphing algorithms; this evaluation is motivated by the fact that the same morphing algorithms have been used for morphing generation in the two datasets so that possible differences that might arise are mainly related to the nature of the data (synthetic vs. real).

The MAP for this experiment has been computed for the two datasets by comparing each morphed image with a single probe image of each contributing subject (the FEI dataset includes only 2 "ISO/ICAO compliant" images, so one of them has been used for morphing generation and the other one as probe image). The MAP matrix contains therefore one single row and three columns corresponding to the number of commercial FRSs used in the evaluation.

The analysis of the results, reported for the different morphing algorithms, show that both datasets present a very high attack potential for all the algorithms analyzed, with C02 and C05 reaching slightly but constantly higher results. This result validates the effectiveness of synthetic data in attacking commercial FRSs to a similar extent than real data. However, in general MONOT MAP is slightly lower than FEI MAP, and a further investigation has been conducted to better understand this phenomenon. In particular, we performed a number of bona fide verification attempts where two images of the same subject are compared and the score distributions in the two datasets obtained with the three FRSs is analyzed. The scores measured are all definitely higher than the FAR 0.1% verification threshold fixed by the FRSs, meaning that the probe images also for synthetic data are correcly recognized as belonging to the same subject. However, the results, reported in Figure 6, clearly show that the bona fide scores for the FEI dataset are higher than those of ONOT (from which MONOT is generated), suggesting that this second dataset is characterized by a higher intra-class variability. This observation is reasonable if we consider that one of the main challenges in synthetic data generation is identity preservation between multiple images of the same virtual subject. We believe that this result explains the slightly lower attack potential observed for MONOT, which in any case achieves performance comparable to the real dataset.

	MONOT			**FEI [57]**		
	1	2	3	1	2	3
C01	**94.2**%	**86.7**%	<u>73.5</u>%	97.9%	91.7%	74.0%
C02	<u>94.1</u>%	**86.7**%	**74.0**%	**99.7**%	**98.9**%	**95.9**%
C03	90.9%	80.8%	63.4%	97.3%	89.9%	70.2%
C05	91.8%	<u>83.0</u>%	67.4%	98.4%	93.5%	78.0%
C15	89.3%	78.3%	55.8%	97.7%	91.5%	70.9%
C16	91.8%	82.3%	61.6%	<u>99.3</u>%	<u>97.4</u>%	<u>88.7</u>%

TABLE I: MAP metrics measured for the MONOT and FEI datasets, for different morphing algorithms (rows). In columns, the number of FRSs considered is reported. Best results in **bold**, second ones <u>underlined</u>.

	1	2	3
1	91.3%	80.1%	60.1%
2	86.7%	73.0%	52.5%
3	83.5%	67.9%	47.1%
4	80.3%	63.4%	42.6%
5	77.3%	59.2%	38.4%
6	74.0%	54.9%	34.5%
7	70.5%	50.5%	30.4%
8	66.6%	45.1%	26.0%
9	61.3%	38.8%	20.6%
10	52.1%	28.0%	12.9%

TABLE II: MAP on the MONOT dataset. In rows, the number of gate probe images generated with Arc2Face [50] is reported, while in columns is the numbers of FRSs considered.

We further analyzed the MONOT attack potential by taking into account the gate images generated using Arc2Face as described in Section III-B. In this case, we have 10 probe gate images for each morphed image for each of the contributing subjects. The MAP matrix for the is reported in Table II. Also in this case, very good results are observed, thus confirming both the validity of the MONOT morphed images and the effectiveness of Arc2Face in preserving the subject's identity across the generation of multiple images.

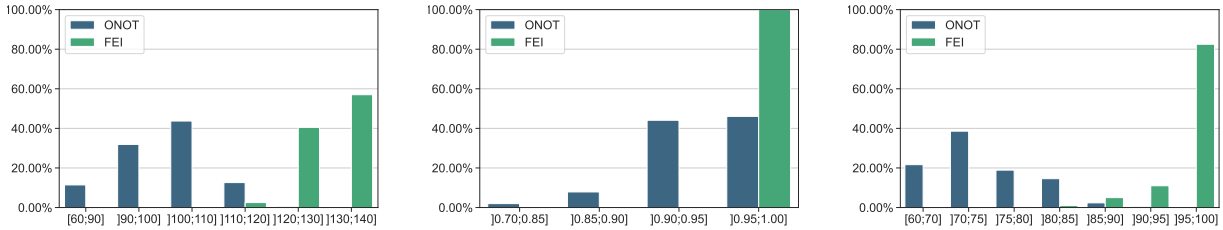

(a) FRS_1 (b) FRS_2 (c) FRS_3

Fig. 6: Bona fide score distributions for the ONOT [43] and FEI [57] datasets, computed with the three FRSs.

	Baseline				Synth				Mix			
	EER	$\mathbf{B}_{0.1}$	$\mathbf{B}_{0.05}$	$\mathbf{B}_{0.01}$	**EER**	$\mathbf{B}_{0.1}$	$\mathbf{B}_{0.05}$	$\mathbf{B}_{0.01}$	**EER**	$\mathbf{B}_{0.1}$	$\mathbf{B}_{0.05}$	$\mathbf{B}_{0.01}$
C02	.2650	.5716	.7040	.9110	.3928	.8171	.9013	.9868	**.1426**	**.2082**	**.3466**	**.6438**
C03	.1795	.3502	.5415	.8436	.4308	.8797	.9495	.9952	**.0893**	**.0722**	**.1745**	**.4753**
C08	**.1847**	.3730	.5788	.8929	.5945	.9061	.9519	.9904	.2265	**.0722**	**.1745**	**.4753**

TABLE III: S-MAD results obtained on the ChiMO dataset [33] for the models trained using real data only (Baseline), Synthetic data only (Synth) and the union of real and synthetic data (Mix).

C. Results on S-MAD training

A second set of experiments has been conducted to evaluate the utility of MONOT synthetic data in S-MAD systems training. For our evaluation, we adopt here the common metrics used for this task, namely the Bona fide Classification Error Rate (BPCER) and the Morphing Attack Classification Error Rate (MACER) [63], representing the percentage of bona fide images wrongly classified as morphed and the percentage of morphed image classified as bona fide, respectively. In particular, the BPCER at fixed MACER values are reported in the results tables: $\mathbf{B}_{0.1}$, $\mathbf{B}_{0.05}$ and $\mathbf{B}_{0.01}$, corresponding respectively to a MACER of 10%, 5% 1%. Moreover, the Equal Error Rate (EER) is also reported.

Given the quite large number of S-MAD solutions in the literature, we have implemented the S-MAD model detailed in [33], chosen for its state-of-the-art results [35]. In our implementation, we utilize the Inception-ResNet architecture [64] as a binary classifier, distinguishing between "morphed" and "bona fide" classes. The model is trained using the SGD optimizer with a momentum of 0.9 and an initial learning rate set to 10^{-3}. Input faces are first detected and cropped using the MTCNN [65] face detector, known for its superior performance in S-MAD tasks [33].

For this experiment, we perform a challenging cross-dataset and cross-morphing algorithm evaluation. In particular, we compare three S-MAD models trained as follows: *Baseline* - real data only: training on PMDB (morphed images generated with the C05 morphing algorithm) and Idiap FaceMorpher (morphing algorithm C01); *Synth* - synthetic data only: training of the portion of MONOT images generated using the same C05 and C01 morphing algorithms; *Mix* - real and synthetic data: union of the previous two training datasets. Following the experimental protocol described in [33], all the models have been tested on the ChiMO dataset, whose morphed images are generated using three morphing algorithms (C02, C03 and C08), not represented in the training set. The results are reported in Table III for the three models. The error rates observed for the model trained only on synthetic data are higher than those obtained for the baseline model, suggesting that the exclusive use of synthetic images for model training is not feasible at this stage. Conversely, the joint use of real and synthetic data (*Mix*) provides very interesting results, and a significant performance improvement can be observed with respect to the use of real data only (*Baseline*), confirming that synthetic data can be successfully employed to extend the set of training images for the S-MAD task.

V. Discussion about synthetic data

We recognize that using diffusion models trained on extensive web-scraped datasets raises crucial ethical, legal, and privacy concerns. Our aim with the MONOT dataset is to explore alternative face data generation methods that could potentially decrease the dependency on real, identifiable individuals. We acknowledge that this approach may not fully address the complex issues surrounding ethics and privacy in face recognition and related tasks.

VI. Conclusions

This paper introduces the MONOT dataset, which mimics real ISO/ICAO-compliant biometric data without violating privacy, and proves its effectiveness for robust S-MAD model training. The experiments in terms of attack potential, where MONOT achieved results comparable to those of a real dataset, confirm the high quality of the synthetic morphed images. Furthermore, MAD systems trained with the addition of this synthetic dataset exhibit strong detection performance across a challenging cross-dataset and cross-algorithm evaluation. These findings highlight the utility of synthetic data in developing robust MAD systems, partially addressing the ethical and privacy challenges of using real biometric data.

REFERENCES

[1] M. Ferrara, A. Franco, and D. Maltoni, "The magic passport," in *IEEE International Joint Conference on Biometrics, Clearwater, IJCB 2014, FL, USA, September 29 - October 2, 2014*, 2014.

[2] U. Scherhag, C. Rathgeb, J. Merkle, R. Breithaupt, and C. Busch, "Face recognition systems under morphing attacks: A survey," *IEEE Access*, vol. 7, pp. 23 012–23 026, 2019.

[3] K. B. Raja *et al.*, "Morphing attack detection-database, evaluation platform, and benchmarking," *IEEE Trans. Inf. Forensics Secur.*, vol. 16, pp. 4336–4351, 2021.

[4] G. D. P. R. GDPR, "General data protection regulation," *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC*, 2016.

[5] N. C. Abay, Y. Zhou, M. Kantarcioglu, B. Thuraisingham, and L. Sweeney, "Privacy preserving synthetic data release using deep learning," in *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2018, Dublin, Ireland, September 10–14, 2018, Proceedings, Part I 18*. Springer, 2019, pp. 510–526.

[6] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 10 684–10 695.

[7] D. P. Kingma, M. Welling *et al.*, "An introduction to variational autoencoders," *Foundations and Trends® in Machine Learning*, vol. 12, no. 4, pp. 307–392, 2019.

[8] "ISO/IEC 39794-5 — Information technology — Extensible biometric data interchange formats — Part 5: Face image data," International Organization for Standardization, Standard, 2019.

[9] H. Qiu, B. Yu, D. Gong, Z. Li, W. Liu, and D. Tao, "Synface: Face recognition with synthetic data," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10 880–10 890.

[10] G. Bae, M. de La Gorce, T. Baltrušaitis, C. Hewitt, D. Chen, J. Valentin, R. Cipolla, and J. Shen, "Digiface-1m: 1 million digital face images for face recognition," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 3526–3535.

[11] F. Boutros, M. Klemt, M. Fang, A. Kuijper, and N. Damer, "Unsupervised face recognition using unlabeled synthetic data," in *2023 IEEE 17th International Conference on Automatic Face and Gesture Recognition (FG)*. IEEE, 2023, pp. 1–8.

[12] F. Boutros, M. Huber, P. Siebke, T. Rieber, and N. Damer, "Sface: Privacy-friendly and accurate face recognition using synthetic data," in *2022 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2022, pp. 1–11.

[13] L. Colbois, T. de Freitas Pereira, and S. Marcel, "On the use of automatically generated synthetic image datasets for benchmarking face recognition," in *2021 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2021, pp. 1–8.

[14] L. Guarnera, O. Giudice, F. Guarnera, A. Ortis, G. Puglisi, A. Paratore, L. M. Bui, M. Fontani, D. A. Coccomini, R. Caldelli *et al.*, "The face deepfake detection challenge," *Journal of Imaging*, vol. 8, no. 10, p. 263, 2022.

[15] F. Boutros, J. H. Grebe, A. Kuijper, and N. Damer, "Idiff-face: Synthetic-based face recognition through fizzy identity-conditioned diffusion model," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 19 650–19 661.

[16] M. Kansy, A. Raël, G. Mignone, J. Naruniec, C. Schroers, M. Gross, and R. M. Weber, "Controllable inversion of black-box face recognition models via diffusion," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 3167–3177.

[17] P. Dhariwal and A. Nichol, "Diffusion models beat gans on image synthesis," *Advances in neural information processing systems*, vol. 34, pp. 8780–8794, 2021.

[18] A. Q. Nichol and P. Dhariwal, "Improved denoising diffusion probabilistic models," in *International conference on machine learning*. PMLR, 2021, pp. 8162–8171.

[19] A. N. Escalante B and L. Wiskott, "Gender and age estimation from synthetic face images," in *Computational Intelligence for Knowledge-Based Systems Design: 13th International Conference on Information Processing and Management of Uncertainty, IPMU 2010, Dortmund, Germany, June 28-July 2, 2010. Proceedings 13*. Springer, 2010, pp. 240–249.

[20] 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," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 1606–1617.

[21] "ISO/IEC 19794-5 — Information technology — Biometric data interchange formats — Part 5: Face image data," International Organization for Standardization, Standard, 2011.

[22] A. Franco, A. Magnani, D. Maltoni, D. Maio, L. Odorisio, and A. De Maria, "Face image quality assessment in electronic id documents," *IEEE Access*, vol. 10, pp. 77 744–77 758, 2022.

[23] T. Schlett, C. Rathgeb, O. Henniger, J. Galbally, J. Fierrez, and C. Busch, "Face image quality assessment: A literature survey," *ACM Computing Surveys (CSUR)*, vol. 54, no. 10s, pp. 1–49, 2022.

[24] J. M. Singh and R. Ramachandra, "3d face morphing attacks: Generation, vulnerability and detection," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 2023.

[25] 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.

[26] 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.

[27] N. Damer, M. Fang, P. Siebke, J. N. Kolf, M. Huber, and F. Boutros, "Mordiff: Recognition vulnerability and attack detectability of face morphing attacks created by diffusion autoencoders," in *2023 11th International Workshop on Biometrics and Forensics (IWBF)*. IEEE, 2023, pp. 1–6.

[28] M. Ferrara, A. Franco, and D. Maltoni, "Face demorphing," *IEEE Trans. Inf. Forensics Secur.*, vol. 13, no. 4, 2018.

[29] G. Borghi, A. Franco, G. Graffieti, and D. Maltoni, "Automated artifact retouching in morphed images with attention maps," *IEEE Access*, vol. 9, pp. 136 561–136 579, 2021.

[30] N. Di Domenico, G. Borghi, A. Franco, and D. Maltoni, "Face restoration for morphed images retouching," in *2024 12th International Workshop on Biometrics and Forensics (IWBF)*. IEEE, 2024, pp. 1–6.

[31] I. Medvedev, F. Shadmand, and N. Gonçalves, "Mordeephy: Face morphing detection via fused classification," *arXiv preprint arXiv:2208.03110*, 2022.

[32] J. Tapia, D. Schulz, and C. Busch, "Single morphing attack detection using siamese network and few-shot learning," *arXiv preprint arXiv:2206.10969*, 2022.

[33] 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.

[34] M. Long, X. Zhao, L.-B. Zhang, and F. Peng, "Detection of face morphing attacks based on patch-level features and lightweight networks," *Security and Communication Networks*, vol. 2022, no. 1, p. 7460330, 2022.

[35] Biolab, "FVC-onGoing." [Online]. Available: https://biolab.csr.unibo.it/fvcongoing/

[36] R. Ramachandra and G. Li, "Residual colour scale-space gradients for reference-based face morphing attack detection," in *2022 25th International Conference on Information Fusion (FUSION)*. IEEE, 2022, pp. 1–8.

[37] S. Soleymani, A. Dabouei, F. Taherkhani, J. Dawson, and N. M. Nasrabadi, "Mutual information maximization on disentangled representations for differential morph detection," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021, pp. 1731–1741.

[38] S. Soleymani, B. Chaudhary, A. Dabouei, J. Dawson, and N. M. Nasrabadi, "Differential morphed face detection using deep siamese networks," in *International Conference on Pattern Recognition*. Springer, 2021, pp. 560–572.

[39] S. Banerjee and A. Ross, "Conditional identity disentanglement for differential face morph detection," in *2021 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2021, pp. 1–8.

[40] U. Scherhag, L. Debiasi, C. Rathgeb, C. Busch, and A. Uhl, "Detection of face morphing attacks based on PRNU analysis," *IEEE Trans. Biom. Behav. Identity Sci.*, vol. 1, no. 4, 2019.

[41] P. Aghdaie, B. Chaudhary, S. Soleymani, J. Dawson, and N. M. Nasrabadi, "Attention aware wavelet-based detection of morphed face images," in *2021 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2021, pp. 1–8.

[42] M. Long, C.-k. Jia, and F. Peng, "Face morphing detection based on a two-stream network with channel attention and residual of multiple color spaces," in *International Conference on Machine Learning for Cyber Security*. Springer, 2022, pp. 439–454.

[43] N. Di Domenico, G. Borghi, A. Franco, D. Maltoni *et al.*, "Onot: a high-quality icao-compliant synthetic mugshot dataset," in *The 18th IEEE International Conference on Automatic Face and Gesture Recognition (FG)*, 2024, pp. 1–6.

[44] A. Quek, "FaceMorpher morphing algorithm." [Online]. Available: https://github.com/alyssaq/face_morpher

[45] FaceFusion, "Facefusion." [Online]. Available: http://www.wearemoment.com/FaceFusion/

[46] M. Ferrara, A. Franco, and D. Maltoni, "Decoupling texture blending and shape warping in face morphing," in *International Conference of the Biometrics Special Interest Group (BIOSIG)*, September 2019.

[47] SURYS Group, "SURYS web site." [Online]. Available: https://surys.com/

[48] INGROUPE, "INGROUPE web site." [Online]. Available: https://ingroupe.com/

[49] I. Batskos, L. Spreeuwers, and R. Veldhuis, "Visualizing landmark based face morphing traces on digital images," *Frontiers in Computer Science*, vol. 5, 2023.

[50] F. P. Papantoniou, A. Lattas, S. Moschoglou, J. Deng, B. Kainz, and S. Zafeiriou, "Arc2face: A foundation model of human faces," *arXiv preprint arXiv:2403.11641*, 2024.

[51] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, 2019.

[52] A. Martinez and R. Benavente, "The AR face database: Cvc technical report, 24," 1998.

[53] P. J. Phillips *et al.*, "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.

[54] P. J. Phillips, H. Wechsler, J. Huang, and P. J. Rauss, "The FERET database and evaluation procedure for face-recognition algorithms," *Image and vision computing*, vol. 16, no. 5, 1998.

[55] E. Sarkar, P. Korshunov, L. Colbois, and S. Marcel, "Vulnerability analysis of face morphing attacks from landmarks and generative adversarial networks," *arXiv preprint arXiv:2012.05344*, 2020.

[56] E. Sarkar, P. Korshunov, L. Colbois, and M. Sebastien, "Are gan-based morphs threatening face recognition?" in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 2959–2963.

[57] 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.

[58] C. Thomaz and G. Giraldi, "A new ranking method for principal components analysis and its application to face image analysis," *Image and Vision Comput.*, 2010.

[59] K. Raja *et al.*, "Morphing attack detection-database, evaluation platform, and benchmarking," *IEEE transactions on information forensics and security*, vol. 16, pp. 4336–4351, 2020.

[60] D. S. Ma, J. Correll, and B. Wittenbrink, "The chicago face database: A free stimulus set of faces and norming data," *Behavior research methods*, vol. 47, no. 4, pp. 1122–1135, 2015.

[61] M. Ferrara, A. Franco, D. Maltoni, and C. Busch, "Morphing attack potential," in *2022 International Workshop on Biometrics and Forensics (IWBF)*, 2022, pp. 1–6.

[62] NIST, "Face recognition vendor test (FRVT) 1:1 verification." [Online]. Available: https://pages.nist.gov/frvt/html/frvt11.html

[63] "ISO/IEC CD 20059.2 Methodologies to evaluate the resistance of biometric recognition systems to morphing attacks," International Organization for Standardization, Standard, 2023.

[64] 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.

[65] 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.