

Privacy-friendly Synthetic Data for the Development of Face Morphing Attack Detectors

Naser Damer^{1,2}, César Augusto Fontanillo López³, Meiling Fang^{1,2}, Noémie Spiller¹
Minh Vu Pham¹, Fadi Boutros^{1,2}

¹Fraunhofer Institute for Computer Graphics Research IGD, Darmstadt, Germany

²Department of Computer Science, TU Darmstadt, Darmstadt, Germany

³Centre for IT & IP Law, KU Leuven, Leuven, Belgium

Email: naser.damer@igd.fraunhofer.de

Abstract

The main question this work aims at answering is: "can morphing attack detection (MAD) solutions be successfully developed based on synthetic data?". Towards that, this work introduces the first synthetic-based MAD development dataset, namely the Synthetic Morphing Attack Detection Development dataset (SMDD). This dataset is utilized successfully to train three MAD backbones where it proved to lead to high MAD performance, even on completely unknown attack types. Additionally, an essential aspect of this work is the detailed legal analyses of the challenges of using and sharing real biometric data, rendering our proposed SMDD dataset extremely essential. The SMDD dataset, consisting of 30,000 attack and 50,000 bona fide samples, is publicly available for research purposes ¹.

1. Introduction

Face recognition (FR) systems, despite their high accuracy [22] and acceptability [3], are vulnerable to many attacks [12, 18, 47], one of these is the face morphing attack (MA) [27]. Ferrara et al. [27] analyzed face MAs by showing that an attack face image can be, automatically and by human experts, matched to more than one person. If MAs targeted travel or identity documents, they would enable multiple subjects to be verified to one document. This faulty link between a document and an identity can lead to illegal activities, including human trafficking, illegal immigration, and questionable financial transactions.

Such attacks can, to a certain degree, be detected using MAD solutions [50]. MAD solutions typically require to be trained to differentiate between the two classes, MAs and bona fide (BF) (i.e. not attack) face images. This train-

ing requires data containing samples of these classes to be used and possibly shared for MAD research and development by different persons and institutions. Such data is, till now, based on real biometric data, limited in diversity and quantity, and raises ethical and legal challenges on its use and distribution, as will be discussed in detail in this paper. This motivates our pursuit of using synthetic-based data for the development of MADs. The pursue is based on the fact that MAD solutions are trained to detect the changes that the morphing process introduces to the image and makes the MA "different" in comparison to BF samples. This detectable difference, therefore, might be detectable as well between synthetic images before and after undergoing a morphing process.

To address the need for large-scale MAD development datasets along with the legal and ethical challenges of using/reusing/sharing real biometric data, we raise the question of "can MAD solutions be successfully developed based on synthetic data?". As a first step towards our goal, we provide detailed and comprehensive legal analyses on the challenges raised by using real biometric data, thus motivating the move towards synthetic data in biometric solutions development. We then present in this work our publicly available and synthetic-based SMDD dataset containing 30,000 MA and 50,000 BF samples. By extensive experiments on three diverse MAD backbones and the evaluation over different MA variations, we can state that the MADs trained on our synthetic SMDD data did lead to stable MAD performances even on MAs created by unknown (not included in the training) morphing techniques.

2. Related work

The creation of face MA datasets started with the initial works that uncovered the attack possibility [27]. This was followed by a set of datasets that were created to train and evaluate MADs [53–55] and to further study the vulnerability of FR to MAs [28, 59]. Most of such datasets cre-

¹<https://github.com/naserdamer/SMDD-Synthetic-Face-Morphing-Attack-Detection-Development-dataset>

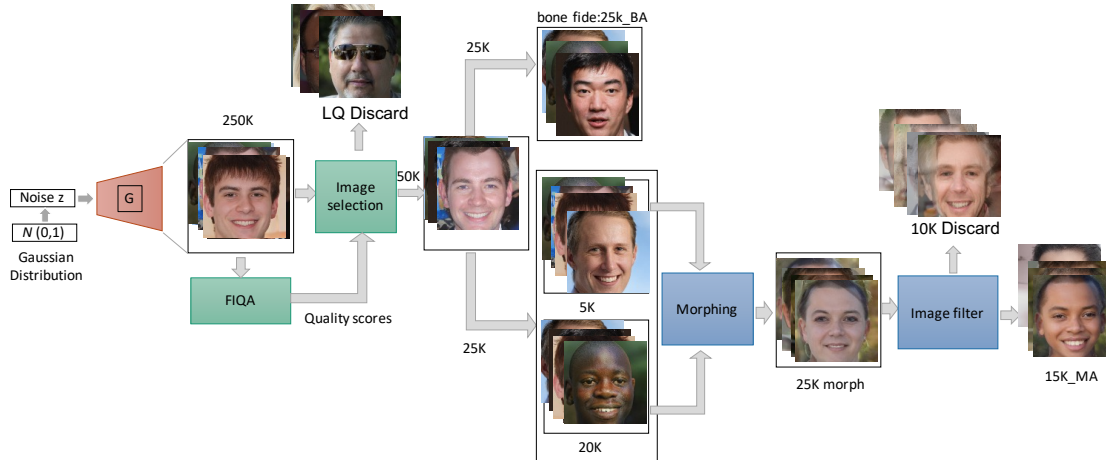


Figure 1. Workflow of the SMDD dataset creation, starting from synthetic faces to the final MA and BF sets.

ated MAs based on the interpolation of facial landmarks, whether detected manually or automatically [28, 53, 60]. Another face morphing strategy based on generative adversarial networks (GAN) was proposed in [15]. This was later enhanced with a cascaded refinement network to possess higher perceived quality [11]. Using more advanced GAN architectures, [85] and [83] produced more realistic GAN generated morphs. The post-processing of landmark-based face morphs in a GAN architecture was later proposed in [14]. Some of the created datasets focused on creating versions with certain variations such as morphing pair selection protocols [16, 19], print and scan operations [17, 31], or variations in landmark based morphing [57, 58, 60].

Most of these existing datasets only included up to a thousand MAs per morphing strategy [10, 11, 13, 15, 17, 27–31, 34, 53, 55, 57–59], which might be insufficient to train an over-parameterized deep neural network and had very limited variation to create a generalizable MAD. Larger datasets were created for this reason with over a thousand morphs per morphing strategy, but also were restricted to under 4 thousand morphs [14, 54, 56, 60, 83, 85], which is still hundreds of folds smaller in comparisons to datasets used for common FR [35, 84] or spoof detection training [87].

From these discussed MA datasets, only a few are available for MAD development research. To the best of our knowledge, the available datasets are the MorGan and associated LMA sets with 1000 MAs, and their enhanced versions [11, 15], the limited print and scan LMA-DRD dataset with only 276 attacks [17], and the recently released FRLM-Morphs, FERET-Morphs, and FRGC-Morphs, each with under or around one thousand attacks per most morph types [57]. FRLM-Morphs, FERET-Morphs, and FRGC-Morphs datasets are not presented as an MAD training or evaluation set, and they aim at enabling the evaluation of FR vulnerability to MAs.

The fact of the limited size of the existing face MA

datasets, the privacy issues that limits their public availability, and the ethical concerns on sharing and reusing biometric information of individuals, stresses the need for MAD development datasets that are of a large-scale, privacy-friendly, and can be shared in the research community. This gap is explored here by the proposal of the SMDD dataset.

3. The legal need for synthetic-based solutions

Biometric data processing is a legally challenging undertaking. Since the entry into force of the General Data Protection Regulation (GDPR) [82], biometric data has acquired legal recognition of a special category of personal data while being subject to a strict data protection regime [76]. According to this body of law, the processing of special categories of personal data, including biometric data, requires the highest standards of protection in view of the fundamental rights and freedoms of individuals. This translates into onerous legal obligations for biometric data controllers, including, among others, (i) the necessity to adhere to one of the exemptions of biometric data processing cumulatively with the choice of a suitable legal basis [77]; (ii) the need to comply with the national law of the European Union (EU) Member State relating to the processing of biometric data, where applicable [81]; (iii) the maintenance of processing records [67]; and (iv) the preparation of a data protection impact assessment and the appointment of a DPO, where applicable [68, 69]. Moreover, depending on the circumstances and the intended purposes of the biometric data processing, data controllers would also be required to (v) comply with the mandatory obligations relating to automated individual decision-making, where applicable [64]; (vi) assume the loss of any possibility of exemption with respect to the appointment of a representative in the EU for controllers and processors located outside the EU [66], and, (vii) where biometric data is processed for purposes other than those initially foreseen, consider the merits and demerits of their compatibility [74]. These legal obligations are,

in common practice, supplemented by other ethical requirements, such as the approval of the biometric data processing by the respective ethics committee or competent authority prior to its onset, particularly in the academic sector ².

Given the exigent legal and ethical framework, the carrying out of biometric data processing can become a tedious, time-consuming, and resource-intensive task. More particularly, biometric data sharing and reuse are deemed as some of the most challenging situations faced by data controllers due to the bureaucracy imposed by data protection law. Compliance with the GDPR would require, in these cases, the approach to the data subjects to request their explicit consent where no other exception and legal basis can provide for a lawful and fair sharing or reuse of the biometric data, and where the 'compatibility test' [74] cannot offer a legally viable solution. For instance, where the reuse of biometric data cannot be based on the scientific research exception [80] coupled with the need to fulfill a task performed for reasons of public interest [73], and the 'compatibility test' does not offer a positive result in light of, inter alia, the sensitivity of biometric data [75], the principle of fairness [70] would require, considering the inadequacy of the other available exceptions [79] and legal bases [72], recourse to explicit consent [78] cumulatively with the consent [71] of the data subject as the appropriate exception and legal basis, respectively, for the reuse of biometric data. Furthermore, the exercise of data subject's rights may, in some instances, jeopardize these practices or even be unrealizable for the biometric data controller ³. In this vein, compliance with the withdrawal of consent may, depending on the number of requests, allegedly undermine the subsequent reuse of biometric data. This may, in turn, impact the high investments that public and private institutions make for the furtherance of biometrics. Similarly, the exercise of the right of erasure by the data subject after the sharing or publication of its biometric data may not be possible or may not be effectively enforced by the data controller once the biometric data have left its sphere of control. In a global context of competition and European cohesion characterized by the pressing need to institute growing databases rich in biometric data to advance the state-of-the-art in scientific and industrial applications as well as data governance [7–9, 32, 51], it is debatable how the free flow of information can be achieved in the EU while ensuring strict compliance with the right to the protection of personal data. More importantly, it may be questionable how public

and private actors may comply with the stringent regime applicable to biometric data without compromising or seriously impairing the achievement of biometric data sharing or reuse. Based on these assumptions, the legally admissible regime for the processing of biometric data could end up being a practically unbearable regime for data controllers, primarily where databases comprise the biometric data of a large number of data subjects.

To overcome the dilemma of 'protection of personal data - free flow of information', data controllers recurrently seek to anonymize personal data. However, since the threshold for anonymization is being set very high by the European legislator [65], notorious expenditures in terms of data transformation procedures and control mechanisms must be made by the data controllers to anonymize personal data. In addition to this, the anonymization of personal data creates other undesired situations, such as the trade-off between privacy and utility [48]. It is for these reasons that more effective data-driven approaches have been recently explored to cope with the pitfalls of current anonymization techniques [23, 43, 45]. Among them, the use of synthetic data is gaining increasing attention in the technical and legal community due to its towering advantages over conventional anonymization techniques [24]. It is argued by some that the particular properties of synthetic data will offer an alternative to private and public actors for research and knowledge transfer, while others have doubts [1]. Advantages of synthetic data include that data controllers may circumnavigate the legal constraints of biometric data processing while maintaining utility, thus permitting the alignment of innovation and law [25], the main motivation behind this work.




Figure 2. Samples of the synthetic SMDD dataset. The top and bottom rows are BF samples. The middle row is samples of the attacks created from the corresponding images above and below the morphed image.

²For instance, copies of opinions and/or approvals by ethics committees and/or competent authorities for biometric data collections are common requirements under EU-funded projects. In certain cases, ethics advisors are also recruited to ensure the responsible, morally acceptable and legally compliant research oversight. E.g. image Manipulation Attack Resolving Solutions (iMARS) Project, <https://imars-project.eu/>

³Most prominently, the right to withdraw the consent of Art. 7(3) GDPR [63] and the right to erasure of Art. 17 GDPR [62].

4. The SMDD dataset

This section presents our Synthetic Morphing Attack Detection Development dataset (SMDD). This is the first synthetic-based biometric attack detection dataset to be proposed and released to the research community. The overview of the dataset creation process is presented in Fig. 1. To start, we created 500k synthetic images using the official open source implementation⁴ of StyleGAN2-ADA [38] trained on Flickr-Faces-HQ Dataset (FFHQ) dataset [39]. These images are randomly generated based on a Gaussian noise Z , drawn without repetition 500k times, from a normal distribution. From each latent vector Z , the pretrained generator produces a synthetic face data sample. These data samples are split randomly into 250k samples that will be the basis of the training data, and a 250k samples that will be the basis of the evaluation data. The main goal of the SMDD dataset is to provide a training dataset for MADs, however, a disjoint testing dataset is created to investigate the realistic deployability of such synthetic datasets as an evaluation benchmark as well. For each of these two sets, the same process described in Fig. 1 is performed. The 250k samples are filtered automatically by using the 50k samples with the highest face image quality, calculated by the CR-FIQA (S) quality assessment approach [5]. The CR-FIQA (S) quality assessment approach [5] reflects the utility of the sample for FR, as defined in ISO/IEC 29794-1 [37], and thus reflects extreme issues that effect FR performance (e.g. extreme poses or occlusions) in a low quality score. This automatic filtering removed most of the images with extreme non-frontal poses and largely occluded (sunglasses and generation artifacts). These filtered 50k images are randomly split into two 25k parts, the first will be considered the BF samples, namely the 25k-BF, and will not go into a morphing process. The second part will be used to create the face MAs. 5k of these images are chosen randomly as the key morphing images, each is paired with 5 images chosen randomly from the 20k left images of the 25k images aimed to create the attacks. This results in 25k morphing pairs. The morphing pairs are selected randomly to allow for a larger variety of similarities between the pairs. Although high similarity pairs create stronger attacks (more successful at attacking FR systems) [16], having variation in the pair similarity in the MAD training data has shown to lead to more generalizable MADs [16]. Given that training MADs is the main reason behind the creation of SMDD, we chose the random selection of morphing pairs. Each of the pairs is morphed using the widely used [15, 57, 60] OpenCV/dlib morphing algorithm [46] based on the “Face Morph Using OpenCV” tutorial⁵ using the Dlib [42] implementation of [40] landmark detector. This morphing tool

⁴<https://github.com/NVlabs/stylegan2-ada>
⁵<http://www.learnopencv.com/face-morph-using-opencv-cpp-python/>




Figure 3. BF samples from the (a) LMA-DRD (D) [17], (b) LMA-DRD (PS) [17], (c) MorGAN-LMA/GAN [15] (MorGan-LMA and MorGan-GAN share bona fide sets), (d) FRLL-Morphs [57].

proved to produce MAs stronger than many other comparable landmark-based morphing tools [57]. The facial landmarks are detected in each of the morphed images [40]. Then, a Delaunay triangulation is performed on each image based on these detected landmarks. Based on calculated affine transforms, the triangles are wrapped with a wrapping factor and the images are blended with a blending factor (also known as a morphing factor). We follow the recommended factor, leading to the most successful attack setups in [30] by setting the blending factor to 0.5 and choosing the warping factor randomly from the range [0.0,1.0]. Examples of the morphed images in our SMDD, along with the synthetic images that created the morphs, are presented in Fig. 2, where one can note the realistic appearance of the MA and BF samples. The resulting 25k MA images are filtered again by removing images with common automatic landmark-based morphing artifacts (black regions on the mouth area) and by manually inspecting the morphed images and disregarding ones with very strong morphing artifacts. This results in reducing the 25k morphed images down to the final set of 15k MAs, namely the 15k-MA. The same process is duplicated for the disjoint training and evaluation sets resulting in a training set of 25k BF and 15k MA samples (the T-25k-BF and T-15k-MA), and the evaluation set of 25k BF and 15k MA samples (the E-25k-BF and E-15k-MA). Examples of the BF and MA samples in the SMDD dataset is presented in Fig. 2.

The complete dataset is available publicly for research purposes. One limitations, and opportunity, of the SMDD dataset is that it only creates MAs using one morphing strategy. However, to enable extending such privacy-friendly dataset in the research community, both the images used to create the MAs and the exact morphing pairs list will be publicly available, where researchers can use them to build subsequent synthetic-based face morphing datasets.

5. MAD and experimental setup

To explore the deployability of the SMDD dataset as a development set for face MAD, we train three different MAD backbones using the SMDD training set. We eval-




Figure 4. Morphing attack samples from the (a) LMA-DRD (D) [17], (b) LMA-DRD (PS) [17], (c) MorGAN-GAN [15], (d) MorGAN-LMA [15], and (e-i) the different attack types of the FRLL-Morphs [57].

uate the resulting MAD on a set of different MA datasets with variations including print and scan operations and a wide range of morphing techniques, including landmark-based and GAN-based ones. To compare the applicability of SMDD as an MAD development set, we compare the performance of the MADs trained on SMDD with the same set of MADs trained on four other datasets. This comparison is not direct as the synthetic nature of the SMDD dataset allows to create a much larger number of samples, unlike conventional real data, however, the performance comparison aims at putting the achieved performances in perspective. This section introduces technical details about the used MADs, the baseline training datasets, the test datasets, the experimental settings, and the MAD evaluation metrics.

5.1. Morphing attack detectors

To give a generalized look at the possibility of using SMDD for MAD, we use three diverse MAD backbones.

Inception-MAD: This baseline uses the Inception-v3 [61] network architecture as its backbone. This architecture has been used successfully for MAD [17, 56] and fake face detection approaches [41]. We report the results of a trained from scratch Inception-v3 model, named InceptionMAD. In the training phase, the binary cross-entropy loss function and Adam optimizer with a learning rate of 10^{-4} and a weight decay of 10^{-5} are used. The input of the network is of the size $229 \times 229 \times 3$. Moreover, early stopping technique with the patience of 20 epochs is used to avoid overfitting and for a fair comparison.

MixFaceNet-MAD: The second utilized backbone architecture is MixFaceNet [4], which is an extremely efficient architecture for FR. MixFaceNet possesses lower computation complexity (FLOPs) and high accuracy and was used successfully in face presentation attack detection [26]. Moreover, MixFaceNet architecture takes advantage of different sizes of convolutional kernels to capture different levels of face attack clues [26]. In our experiment, MixFaceNet, specifically MixFaceNet-S, is used as it represent extremely efficient network. The input image size [4] of MixFaceNet is modified to $224 \times 224 \times 3$ and a fully-connected layer is added after the embedding layer to per-

form the MAD decision. In the training phase, binary cross-entropy loss function, Adam optimizer, and early stopping are used with the same hyper-parameters with the training of Inception-MAD.

PW-MAD: As a complement to the binary label supervision, PW-MAD [17] utilizes pixel-wise map supervision [33], i.e. a label of attack or BF for each image pixel. Considering the superior MAD performance obtained by the PW-MAD method [17], we choose the PW-MAD [17] as our third MAD. The loss for training PW-MAD is the sum of the binary cross-entropy loss computed based on pixel-wise feature map and binary output following the definition in [17]. Our solution replaces the last sigmoid layer in [17] with a softmax layer which proved to enhance generalizability over unknown morphing techniques. The input of the network is of the size $224 \times 224 \times 3$. The hyper-parameters of Adam and early stopping technique is same to used in Inception-MAD and MixFaceNet-MAD.

It must be noted that all the MAD solutions used only the training part of the respective datasets for training, which is identity disjoint from all the testing sets. All the images used for the training and testing are cropped using the MTCNN [86] bounding box detection with 5% extension of the width and height to include the whole face following [17] (examples in Fig. 3 and 4) and all crops are resized to 224×224 .

5.2. Datasets

Beside our SMDD, we utilize five other datasets, LMA-DRD (PS) [17], LMA-DRD (D) [17], MorGAN-LMA [15], MorGAN-GAN [15], and FRLL-Morphs [57]. Of these, the first four contain identity-disjoint test and training sets, where the latter are used to train MAD and put the MAD trained on SMDD in perspective. The fifth contains only test data with five different morphing attack techniques and is used to evaluate the resulting MADs on unknown (in terms of morphing technique and original dataset) attacks.

The **LMA-DRD (D)** and **LMA-DRD (PS)** are built on selected images from the VGGFace2 [6] dataset. These are selected to be frontal, of high resolution, and with neutral expression [17]. The morphing method is based on

Test	MAD	Train Data	EER (%)	BPCER (%) @ APCER =			
				0.10%	1.00%	10.00%	20.00%
FRLL-Morphs - OpenCV	MixFacenet	LMA-DRD (D)	8.83	82.84	34.80	8.82	3.43
	MixFacenet	LMA-DRD (PS)	13.22	64.71	46.08	16.18	7.84
	MixFacenet	MorGAN-LMA	8.91	56.86	17.16	8.33	4.41
	MixFacenet	MorGAN-GAN	17.66	72.06	66.67	25.98	16.67
	MixFacenet	SMDD	4.39	64.22	26.47	1.96	1.47
	PW-MAD	LMA-DRD (D)	17.33	90.20	61.27	24.02	14.22
	PW-MAD	LMA-DRD (PS)	15.69	98.04	90.69	22.06	11.76
	PW-MAD	MorGAN-LMA	13.96	68.14	52.94	17.65	10.29
	PW-MAD	MorGAN-GAN	45.59	98.04	92.65	76.47	66.18
	PW-MAD	SMDD	2.42	71.08	22.06	0.49	0.49
	Inception	LMA-DRD (D)	13.72	86.27	74.02	20.10	11.27
	Inception	LMA-DRD (PS)	10.76	88.24	46.57	10.78	5.39
	Inception	MorGAN-LMA	6.86	56.37	29.41	4.90	1.96
	Inception	MorGAN-GAN	55.89	100.00	100.00	98.04	93.63
Inception	SMDD	5.38	71.57	38.73	1.96	0.98	
FRLL-Morphs - FaceMorpher	MixFacenet	LMA-DRD (D)	7.80	49.51	25.49	5.39	2.45
	MixFacenet	LMA-DRD (PS)	10.97	58.82	42.65	11.27	4.90
	MixFacenet	MorGAN-LMA	7.34	58.82	13.73	5.88	1.96
	MixFacenet	MorGAN-GAN	15.65	68.63	52.45	17.65	12.75
	MixFacenet	SMDD	3.87	63.24	23.53	0.49	0.49
	PW-MAD	LMA-DRD (D)	13.88	81.37	63.73	19.61	10.29
	PW-MAD	LMA-DRD (PS)	15.14	95.59	92.16	20.10	10.29
	PW-MAD	MorGAN-LMA	10.92	59.31	35.29	11.27	4.41
	PW-MAD	MorGAN-GAN	44.57	96.57	91.67	76.47	66.18
	PW-MAD	SMDD	2.20	63.73	26.47	0.49	0.00
	Inception	LMA-DRD (D)	16.62	94.12	78.92	29.41	13.24
	Inception	LMA-DRD (PS)	15.81	97.55	61.27	23.04	12.75
	Inception	MorGAN-LMA	6.32	73.04	23.04	3.92	0.98
	Inception	MorGAN-GAN	66.14	100.00	100.00	100.00	99.02
Inception	SMDD	3.17	70.59	30.39	0.49	0.49	
FRLL-Morphs - StyleGAN2	MixFacenet	LMA-DRD (D)	20.07	91.18	71.57	33.33	20.10
	MixFacenet	LMA-DRD (PS)	15.29	94.12	75.00	23.04	11.76
	MixFacenet	MorGAN-LMA	13.41	65.20	52.45	17.16	11.27
	MixFacenet	MorGAN-GAN	23.51	93.14	76.96	38.24	24.02
	MixFacenet	SMDD	8.99	79.90	42.16	8.82	4.41
	PW-MAD	LMA-DRD (D)	29.97	99.02	97.55	61.27	42.16
	PW-MAD	LMA-DRD (PS)	27.64	99.02	94.12	51.47	33.82
	PW-MAD	MorGAN-LMA	18.11	82.84	68.63	33.82	17.65
	PW-MAD	MorGAN-GAN	48.53	99.51	98.04	84.31	73.04
	PW-MAD	SMDD	16.64	100.00	80.39	25.98	13.24
	Inception	LMA-DRD (D)	37.24	98.53	92.65	69.12	55.88
	Inception	LMA-DRD (PS)	19.58	92.16	83.33	31.86	19.61
	Inception	MorGAN-LMA	20.56	77.94	56.86	31.86	20.59
	Inception	MorGAN-GAN	55.03	100.00	100.00	98.04	92.16
Inception	SMDD	11.37	93.63	72.06	13.73	6.86	
FRLL-Morphs - WebMorph	MixFacenet	LMA-DRD (D)	25.97	93.14	83.33	48.53	30.39
	MixFacenet	LMA-DRD (PS)	29.04	93.14	83.82	53.92	40.69
	MixFacenet	MorGAN-LMA	20.61	81.86	71.57	40.20	22.06
	MixFacenet	MorGAN-GAN	30.39	98.04	93.63	60.29	42.65
	MixFacenet	SMDD	12.35	90.69	80.39	15.20	7.84
	PW-MAD	LMA-DRD (D)	33.78	99.02	89.71	65.20	46.08
	PW-MAD	LMA-DRD (PS)	28.51	98.04	89.71	50.00	35.29
	PW-MAD	MorGAN-LMA	35.75	90.20	83.82	62.75	52.45
	PW-MAD	MorGAN-GAN	52.43	98.04	93.63	84.31	75.00
	PW-MAD	SMDD	16.65	91.18	80.39	24.02	13.24
	Inception	LMA-DRD (D)	57.38	100.00	99.51	96.08	86.76
	Inception	LMA-DRD (PS)	58.32	99.51	99.51	94.12	88.24
	Inception	MorGAN-LMA	30.88	91.18	85.29	54.90	39.22
	Inception	MorGAN-GAN	77.42	100.00	100.00	100.00	100.00
Inception	SMDD	9.86	74.51	53.92	9.80	2.94	
FRLL-Morphs - AMSL	MixFacenet	LMA-DRD (D)	24.53	95.10	85.29	49.51	29.90
	MixFacenet	LMA-DRD (PS)	27.59	95.59	87.75	59.80	37.75
	MixFacenet	MorGAN-LMA	19.24	89.71	69.12	36.76	19.12
	MixFacenet	MorGAN-GAN	30.03	98.04	95.10	60.29	42.65
	MixFacenet	SMDD	15.18	76.47	49.51	21.08	11.76
	PW-MAD	LMA-DRD (D)	36.25	99.02	94.12	71.57	53.43
	PW-MAD	LMA-DRD (PS)	32.95	97.55	85.29	58.33	44.61
	PW-MAD	MorGAN-LMA	34.38	94.12	88.24	68.63	56.37
	PW-MAD	MorGAN-GAN	48.52	99.02	98.04	84.31	75.00
	PW-MAD	SMDD	15.18	100.00	96.57	24.02	5.88
	Inception	LMA-DRD (D)	49.02	100.00	98.53	92.16	81.86
	Inception	LMA-DRD (PS)	61.44	99.51	99.51	97.55	91.67
	Inception	MorGAN-LMA	9.80	58.82	33.33	9.80	4.41
	Inception	MorGAN-GAN	86.49	100.00	100.00	100.00	100.00
Inception	SMDD	10.79	83.33	72.06	12.75	4.90	

Table 1. MAD performances on the different attacks of the FRLL-Morphs when each of the three MADs is trained on LMA-DRD (D), LMA-DRD (PS), MorGAN-LMA, MorGAN-GAN, or SMDD training data. The lowest error rates per MAD are in bold.

OpenCV morphing [46] and following parametrization in [54]. BF images are also selected such that it follows the mentioned manual quality check. The dataset is split into two identical sets, one is digital and noted as LMA-DRD (D) and one is re-digitized (print and scan) and noted as LMA-DRD (PS). From both, we used the train and development sets for training (181 morphs and 241 BF images) and the identity-disjoint test set for testing (88 morphs and

Test	MAD	Train Data	EER (%)	BPCER (%) @ APCER =				
				0.10%	1.00%	10.00%	20.00%	20.00%
LMA-DRD (D)	MixFacenet	LMA-DRD (D)	15.68	82.11	69.92	30.08	13.01	
	MixFacenet	LMA-DRD (PS)	18.03	65.04	56.10	30.89	17.89	
	MixFacenet	MorGAN-LMA	17.06	56.91	56.91	32.52	16.26	
	MixFacenet	MorGAN-GAN	25.10	76.42	74.80	42.28	29.27	
	MixFacenet	SMDD	19.42	86.18	79.67	31.71	19.51	
	PW-MAD	LMA-DRD (D)	20.80	71.54	64.23	48.78	21.14	
	PW-MAD	LMA-DRD (PS)	25.10	92.68	77.24	46.34	26.83	
	PW-MAD	MorGAN-LMA	22.34	82.11	79.67	43.90	22.76	
	PW-MAD	MorGAN-GAN	40.21	98.37	91.87	86.18	69.11	
	PW-MAD	SMDD	17.06	91.87	80.49	25.20	12.20	
	Inception	LMA-DRD (D)	7.64	71.54	15.45	3.25	2.44	
	Inception	LMA-DRD (PS)	17.06	69.11	68.29	23.58	17.07	
	Inception	MorGAN-LMA	15.68	65.04	53.66	31.71	9.76	
	Inception	MorGAN-GAN	50.77	99.19	98.37	94.31	76.42	
Inception	SMDD	15.11	81.30	69.92	30.89	12.20		
LMA-DRD (PS)	MixFacenet	LMA-DRD (D)	21.77	80.49	77.24	34.96	24.39	
	MixFacenet	LMA-DRD (PS)	18.44	67.48	55.28	28.46	18.70	
	MixFacenet	MorGAN-LMA	27.05	74.80	74.80	43.90	29.27	
	MixFacenet	MorGAN-GAN	27.05	86.99	63.41	45.53	37.40	
	MixFacenet	SMDD	23.72	93.50	74.80	52.03	24.39	
	PW-MAD	LMA-DRD (D)	26.48	71.54	69.11	46.34	30.89	
	PW-MAD	LMA-DRD (PS)	23.72	74.80	69.11	30.08	25.20	
	PW-MAD	MorGAN-LMA	29.41	89.43	78.86	62.60	47.15	
	PW-MAD	MorGAN-GAN	44.11	100.00	100.00	86.18	68.29	
	PW-MAD	SMDD	20.39	83.74	82.11	34.15	19.51	
	Inception	LMA-DRD (D)	11.37	56.91	39.02	11.38	8.13	
	Inception	LMA-DRD (PS)	12.75	54.47	37.40	13.01	8.13	
	Inception	MorGAN-LMA	22.34	86.99	69.11	43.09	25.20	
	Inception	MorGAN-GAN	38.42	98.37	97.56	80.49	56.91	
Inception	SMDD	19.01	87.80	86.99	39.84	16.26		
MorGAN-LMA	MixFacenet	LMA-DRD (D)	39.42	92.50	88.76	70.95	55.96	
	MixFacenet	LMA-DRD (PS)	22.89	97.59	70.41	36.55	25.57	
	MixFacenet	MorGAN-LMA	10.61	74.97	46.85	11.11	4.42	
	MixFacenet	MorGAN-GAN	46.42	95.05	93.57	82.46	72.56	
	MixFacenet	SMDD	30.12	89.16	82.06	57.83	39.09	
	PW-MAD	LMA-DRD (D)	34.20	94.38	82.06	61.98	46.72	
	PW-MAD	LMA-DRD (PS)	34.14	95.98	92.64	72.29	51.54	
	PW-MAD	MorGAN-LMA	9.71	81.79	39.89	9.77	4.55	
	PW-MAD	MorGAN-GAN	34.37	93.98	90.50	66.27	49.80	
	PW-MAD	SMDD	27.31	88.89	78.85	50.74	34.14	
	Inception	LMA-DRD (D)	38.55	99.60	95.31	69.61	56.36	
	Inception	LMA-DRD (PS)	31.73	99.46	85.27	61.98	42.44	
	Inception	MorGAN-LMA	8.43	53.55	34.00	7.90	4.42	
	Inception	MorGAN-GAN	40.16	99.87	95.58	78.85	65.46	
Inception	SMDD	28.51	92.10	84.20	54.62	37.08		
MorGAN-GAN	MixFacenet	LMA-DRD (D)	53.01	99.20	96.79	87.68	78.45	
	MixFacenet	LMA-DRD (PS)	50.44	99.60	95.72	80.99	72.29	
	MixFacenet	MorGAN-LMA	42.57	99.33	96.92	81.26	67.34	
	MixFacenet	MorGAN-GAN	24.90	99.33	97.99	57.43	31.59	
	MixFacenet	SMDD	42.64	99.73	98.93	81.39	71.62	
	PW-MAD	LMA-DRD (D)	52.04	98.26	96.79	87.68	79.12	
	PW-MAD	LMA-DRD (PS)	46.59	100.00	98.13	86.21	75.37	
	PW-MAD	MorGAN-LMA	42.80	98.80	93.17	78.45	60.64	
	PW-MAD	MorGAN-GAN	8.84	94.11	40.96	7.76	2.68	
	PW-MAD	SMDD	43.78	99.87	99.60	86.21	72.16	
	Inception	LMA-DRD (D)	50.84	99.73	99.20	91.83	79.25	
	Inception	LMA-DRD (PS)	38.79	99.46	96.92	82.60	65.46	
	Inception	MorGAN-LMA	27.14	97.72	87.28	52.07	32.26	
	Inception	MorGAN-GAN	0.40	37.35	0.00	0.00	0.00	
Inception	SMDD	44.34	99.46	97.32	84.20	73.23		
SMDD	MixFacenet	LMA-DRD (D)	10.34	59.59	32.30	10.54	5.80	
	MixFacenet	LMA-DRD (PS)	9.26	39.82	24.30	8.87	5.59	
	MixFacenet	MorGAN-LMA	5.11	16.77				

OpenCV landmark-based morphing [46] or on the GAN-based solution presented in [15], resulting in the MorGAN-LMA and MorGAN-GAN datasets, respectively. Each of these two sets contains identity-disjoint training and testing sets. Each of the training and testing sets contains 500 MA and 750 BF samples. One aspect that makes this dataset less realistic, but essential to study, is that the generated faces were relatively of low resolution of 64×64 pixels.

The **FRLL-Morphs** dataset [57] is generated from the publicly available Face Research London Lab dataset [21]. The dataset include 5 different morphing methods, including OpenCV (OCV) [46], FaceMorpher (FM) [52], StyleGAN 2 (SG) [38, 83], and WebMorpher (WM) [20], and AMSL [49]. Each morph method contains 1222 morphed face images generated using only frontal face images with high resolution the dataset also contains 204 BF samples. The FRLL-Morphs is not created or structured to develop MADs [57] and thus does not contain a training set, it is rather developed for the evaluation FR vulnerability to attacks and the MAD performance. Following this, we use this dataset as testing set to evaluate the developed MADs on diverse MAs. Samples of all the used datasets are provided in Fig. 3 and 4 presenting BF and MA, respectively.

5.3. Evaluation metrics

The MAD performance is presented by the Attack Presentation Classification Error Rate (APCER), i.e. the proportion of MA images incorrectly classified as BF samples, and the BF Presentation Classification Error Rate (BPCER), i.e. the proportion of BF images incorrectly classified as MA samples, as defined in the ISO/IEC 30107-3 [36]. To cover different operation points, and to present the comparative results, we report the BPCER at three different fixed APCER values (0.1%, 1.0%, and 10.0%). The equal error rate (EER) is also reported, which is the BPCER or APCER value at the decision threshold where they are equal. To provide a visual evaluation on a wider operation range, we plot receiver operating characteristic (ROC) curves by plotting the APCER on the x-axis and 1-BPCER on the y-axis at different operational points. It must be noted again that the MAD evaluation was performed only on the identity-disjoint test data as described in Section 5.2.

6. Results

The MAD performances on the different MAs of the FRLL-Morphs data achieved by the three MADs trained on the SMDD and the other four training sets are presented in Tab. 1. For all the morphing types in the FRLL-Morphs data, the MADs trained on the SMDD data outperformed the ones trained on other datasets in most of the evaluation settings. Even when the MADs trained on the SMDD did not score the lowest error, they scored a very close error rate to the lowest one. In general, FaceMorpher and OpenCV MAs were more detectable than

other FRLL-Morphs attacks, scoring EER values below 5% when the MADs were trained on the SMDD. While PW-MAD trained on SMDD performed generally better than MicFaceNet-MAD and Inception-MAD on the FaceMorpher and OpenCV MAs, MixFaceNet-MAD performed the best on StyleGAN 2 MAs, and the Inception-MAD performed the best on the WebMorph and AMSL MAs. Generally, when the training data has a similar morphing strategy as the testing data, e.g. LMA-DRD (D) and FRLL-Morphs OpenCV, the MADs expectedly performed relatively well.

The MAD performances on the test sets of the LMA-DRD (D), LMA-DRD (PS), MorGAN-LMA, MorGAN-GAN, and SMDD, achieved by the three MADs trained on the SMDD and the other four training sets are presented in Tab. 2. Generally, even though test and train sets are identity-disjoint, the MADs trained on a particular training data achieved high MAD performance on the test set from the same data source, e.g. when MAD trained on MorGAN-LMA training set is evaluated on MorGAN-LMA test set. When neglecting the unrealistic intra-dataset evaluations, MADs trained on SMDD achieved the best performances (especially with the PW-MAD) on many occasions and were a very close second in many others. Generally, MorGAN-GAN MAs were the most difficult to detect, however, they pose the least concerns as MAs given the low vulnerability of FR to them and the relatively unrealistic appearance as demonstrated in [83]. In the cases in Tab. 2, the three MADs trained on SMDD achieved comparable performances. It must be mentioned that the SMDD is not presented as an attack (as it is based on synthetic data) and thus detecting it is not of interest, other than questioning if it can be representative as a privacy-friendly evaluation set. Having this in mind, the SMDD test set did behave similarly (trends in the error rates when different MADs are evaluated) as testing data, compared to the test sets that included attacks created by the same morphing approach, such as LMA-DRD (D) and MorGAN-LMA, however, with significantly lower MAD error rates on the SMDD test set, which might be related to the randomized morph pairing process.

The results in Tab. 1 and 2 are represented over a detailed APCER range as ROC curves in Fig. 5, 6, 7, respectively for the MixFaceNet-MAD, Inception-MAD, and PW-MAD-MAD. The ROCs approve the discussion driven from Tab. 1 and 2 and assure the applicability of the privacy-friendly and synthetic-based SMDD data as development data for high performing and high generalized MADs.

7. Conclusion

This work is the first to question the possibility of using synthetic face data as the basis of face MAD development data. This was driven by presenting an extensive legal analysis on the privacy and security issues raised by using real biometric data. Towards that, this work presents the pub-

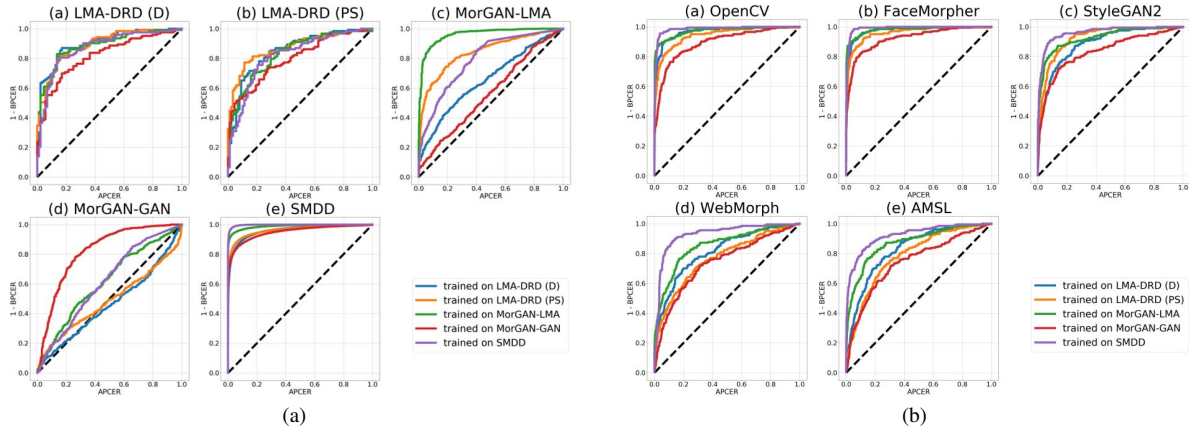


Figure 5. ROC curves achieved on the different test sets and different training settings (data) of the MixFaceNets-MAD. This includes intra-dataset and inter-dataset evaluation in (a) and the different morphing techniques of the FRL-Morphs in (b).

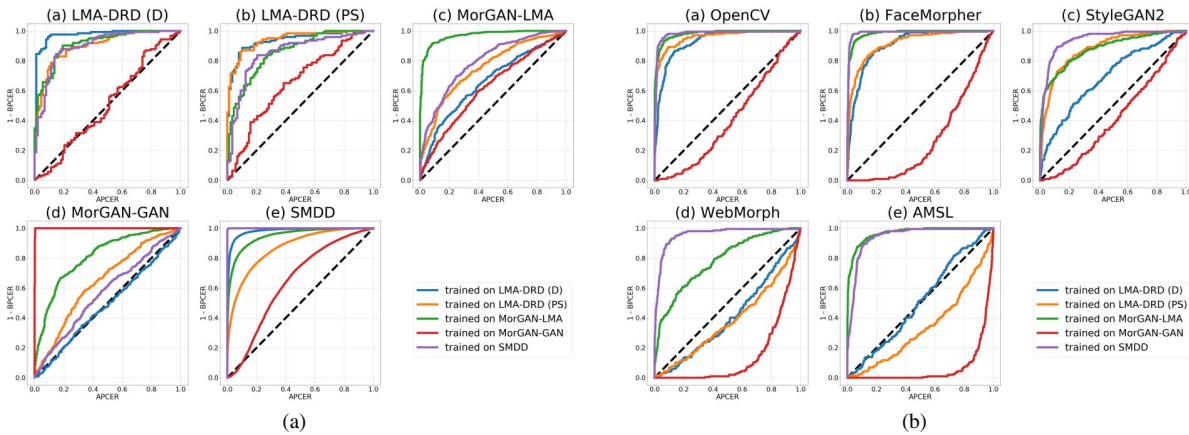


Figure 6. ROC curves achieved on the different test sets and different training settings (data) of the Inception-MAD. This includes intra-dataset and inter-dataset evaluation in (a) and the different morphing techniques of the FRL-Morphs in (b).

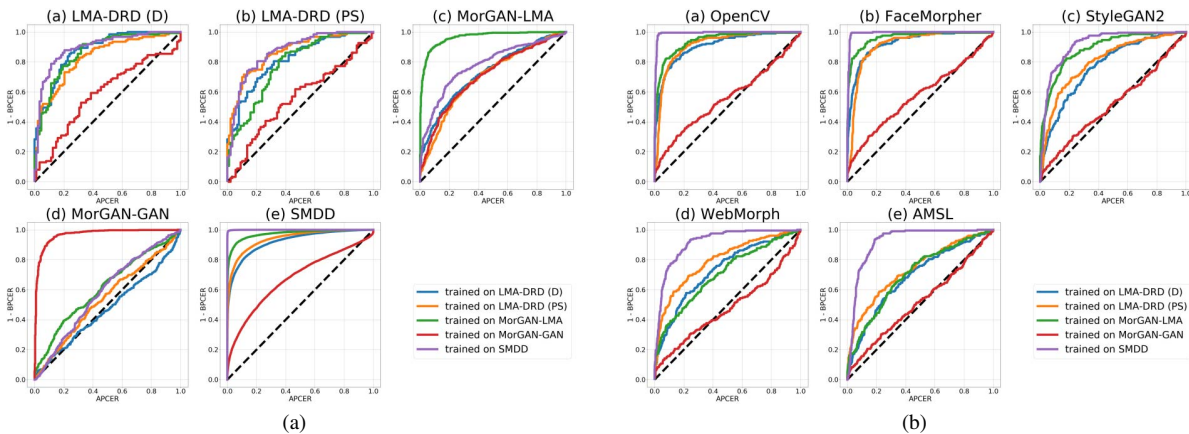


Figure 7. ROC curves achieved on the different test sets and different training settings (data) of the PW-MAD. This includes intra-dataset and inter-dataset evaluation in (a) and the different morphing techniques of the FRL-Morphs in (b).

licly available SMDD dataset containing 30,000 attacks and 50,000 bona fide images, all based on synthetic data. Experiments on three MAD solutions show the applicability of the SMDD as very effective training data, even when the trained MAD is faced with attacks created with unknown morphing techniques or data sources.

Acknowledgment: This research work has been funded by the German Federal Ministry of Education and Research and the Hessen State Ministry for Higher Education, Research and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE.

References

- [1] Synthetic data – anonymisation groundhog day. In *31st USENIX Security Symposium (USENIX Security 22)*, Boston, MA, Aug. 2022. USENIX Association. 3
- [2] Tadas Baltrušaitis, Peter Robinson, and Louis-Philippe Morency. Openface: An open source facial behavior analysis toolkit. In *2016 IEEE Winter Conference on Applications of Computer Vision, WACV 2016, Lake Placid, NY, USA, March 7-10, 2016*, pages 1–10. IEEE Computer Society, 2016. 6
- [3] Ruud Bolle and Sharath Pankanti. *Biometrics, Personal Identification in Networked Society: Personal Identification in Networked Society*. Kluwer Academic Publishers, Norwell, MA, USA, 1998. 1
- [4] Fadi Boutros, Naser Damer, Meiling Fang, Florian Kirchbuchner, and Arjan Kuijper. Mixfacenet: Extremely efficient face recognition networks. In *International IEEE Joint Conference on Biometrics, IJCB 2021, Shenzhen, China, August 4-7, 2021*, pages 1–8. IEEE, 2021. 5
- [5] Fadi Boutros, Meiling Fang, Marcel Klemt, Biying Fu, and Naser Damer. CR-FIQA: face image quality assessment by learning sample relative classifiability. *CoRR*, abs/2112.06592, 2021. 4
- [6] Qiong Cao, Li Shen, Weidi Xie, Omkar M. Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *13th IEEE International Conference on Automatic Face & Gesture Recognition, FG 2018, Xi'an, China, May 15-19, 2018*, pages 67–74. IEEE Computer Society, 2018. 5
- [7] European Commission. White paper on artificial intelligence - a european approach to excellence and trust. 3
- [8] European Commission. Priority societal challenges, 2014. 3
- [9] European Commission. Communication from the commission to the european parliament, the council, the european economic and social committee and the committee of the regions, a european strategy for data, 2020. 3
- [10] Naser Damer, Viola Boller, Yaza Wainakh, Fadi Boutros, Philipp Terhörst, Andreas Braun, and Arjan Kuijper. Detecting face morphing attacks by analyzing the directed distances of facial landmarks shifts. In *Pattern Recognition - 40th German Conference, GCPR 2018, Stuttgart, Germany, October 9-12, 2018, Proceedings*, volume 11269 of *Lecture Notes in Computer Science*, pages 518–534. Springer, 2018. 2
- [11] Naser Damer, Fadi Boutros, Alexandra Mosegui Saladie, Florian Kirchbuchner, and Arjan Kuijper. Realistic dreams: Cascaded enhancement of gan-generated images with an example in face morphing attacks. In *10th IEEE International Conference on Biometrics Theory, Applications and Systems, BTAS 2019, Tampa, FL, USA, September 23-26, 2019*, pages 1–10. IEEE, 2019. 2
- [12] Naser Damer and Kristiyan Dimitrov. Practical view on face presentation attack detection. In *Proceedings of the British Machine Vision Conference 2016, BMVC 2016, York, UK, September 19-22, 2016*. BMVA Press, 2016. 1
- [13] Naser Damer, Jonas Henry Grebe, Steffen Zienert, Florian Kirchbuchner, and Arjan Kuijper. On the generalization of detecting face morphing attacks as anomalies: Novelty vs. outlier detection. In *10th IEEE International Conference on Biometrics Theory, Applications and Systems, BTAS 2019, Tampa, FL, USA, September 23-26, 2019*, pages 1–5. IEEE, 2019. 2
- [14] Naser Damer, Kiran B. Raja, Marius Süßmilch, Sushma Venkatesh, Fadi Boutros, Meiling Fang, Florian Kirchbuchner, Raghavendra Ramachandra, and Arjan Kuijper. Regenmorph: Visibly realistic GAN generated face morphing attacks by attack re-generation. In *Advances in Visual Computing - 16th International Symposium, ISVC 2021, Virtual Event, October 4-6, 2021, Proceedings, Part I*, volume 13017 of *Lecture Notes in Computer Science*, pages 251–264. Springer, 2021. 2
- [15] Naser Damer, Alexandra Mosegui Saladie, Andreas Braun, and Arjan Kuijper. Morgan: Recognition vulnerability and attack detectability of face morphing attacks created by generative adversarial network. In *9th IEEE International Conference on Biometrics Theory, Applications and Systems, BTAS 2018, Redondo Beach, CA, USA, October 22-25, 2018*, pages 1–10. IEEE, 2018. 2, 4, 5, 6, 7
- [16] Naser Damer, Alexandra Mosegui Saladie, Steffen Zienert, Yaza Wainakh, Philipp Terhörst, Florian Kirchbuchner, and Arjan Kuijper. To detect or not to detect: The right faces to morph. In *2019 International Conference on Biometrics, ICB 2019, Crete, Greece, June 4-7, 2019*, pages 1–8. IEEE, 2019. 2, 4
- [17] Naser Damer, Noémie Spiller, Meiling Fang, Fadi Boutros, Florian Kirchbuchner, and Arjan Kuijper. PW-MAD: pixel-wise supervision for generalized face morphing attack detection. In *Advances in Visual Computing - 16th International Symposium, ISVC 2021, Virtual Event, October 4-6, 2021, Proceedings, Part I*, volume 13017 of *Lecture Notes in Computer Science*, pages 291–304. Springer, 2021. 2, 4, 5
- [18] Naser Damer, Yaza Wainakh, Viola Boller, Sven von den Berken, Philipp Terhörst, Andreas Braun, and Arjan Kuijper. Crazyfaces: Unassisted circumvention of watchlist face identification. In *9th IEEE International Conference on Biometrics Theory, Applications and Systems, BTAS 2018, Redondo Beach, CA, USA, October 22-25, 2018*, pages 1–9. IEEE, 2018. 1
- [19] Naser Damer, Steffen Zienert, Yaza Wainakh, Alexandra Mosegui Saladie, Florian Kirchbuchner, and Arjan Kuijper. A multi-detector solution towards an accurate and generalized detection of face morphing attacks. In *22th International Conference on Information Fusion, FUSION 2019, Ottawa, ON, Canada, July 2-5, 2019*, pages 1–8. IEEE, 2019. 2
- [20] Lisa DeBruine. *debruine/webmorph*: Beta release 2. Jan. 2018. 7
- [21] Lisa DeBruine and Benedict Jones. *Face research lab london set*, May 2017. 7
- [22] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos 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*, pages 4690–4699. Computer Vision Foundation / IEEE, 2019. 1

- [23] C. Dwork, M. Bugliese, B. Preneel, V. Sassone, and I. Weinger, editors. *Differential Privacy*, pages 1–12. Springer, Berlin, Heidelberg, 2006. 3
- [24] Khaled El Emam. Seven ways to evaluate the utility of synthetic data. *IEEE Security Privacy*, 18(4):56–59, 2020. 3
- [25] K.E. El Emam, L. Mosquera, and R. Hoptroff. *Practical Synthetic Data Generation: Balancing Privacy and the Broad Availability of Data*. O’Reilly Media, Incorporated, 2020. 3
- [26] Meiling Fang, Fadi Boutros, Arjan Kuijper, and Naser Damer. Partial attack supervision and regional weighted inference for masked face presentation attack detection. In *16th IEEE International Conference on Automatic Face and Gesture Recognition, FG 2021, Jodhpur, India, December 15-18, 2021*, pages 1–8. IEEE, 2021. 5
- [27] Matteo Ferrara, Annalisa Franco, and Davide Maltoni. The magic passport. In *IEEE International Joint Conference on Biometrics, Clearwater, IJCB 2014, FL, USA, September 29 - October 2, 2014*, pages 1–7. IEEE, 2014. 1, 2
- [28] Matteo Ferrara, Annalisa Franco, and Davide Maltoni. On the effects of image alterations on face recognition accuracy. In Thirimachos Bourlai, editor, *Face Recognition Across the Imaging Spectrum*, pages 195–222. Springer, 2016. 1, 2
- [29] Matteo Ferrara, Annalisa Franco, and Davide Maltoni. Face demorphing. *IEEE Trans. Inf. Forensics Secur.*, 13(4):1008–1017, 2018. 2
- [30] Matteo Ferrara, Annalisa Franco, and Davide Maltoni. Decoupling texture blending and shape warping in face morphing. In *2019 International Conference of the Biometrics Special Interest Group, BIOSIG 2019, Darmstadt, Germany, September 18-20, 2019*, volume P-296 of *LNI*, pages 25–33. Gesellschaft für Informatik e.V., 2019. 2, 4
- [31] Matteo Ferrara, Annalisa Franco, and Davide Maltoni. Face morphing detection in the presence of printing/scanning and heterogeneous image sources. *IET Biometrics*, 10(3):290–303, 2021. 2
- [32] Franco Frattini. The commission’s policy priorities in the area of freedom, security and justice, 2 2005. 3
- [33] Anjith George and Sébastien Marcel. Deep pixel-wise binary supervision for face presentation attack detection. In *2019 International Conference on Biometrics, ICB 2019, Crete, Greece, June 4-7, 2019*, pages 1–8. IEEE, 2019. 5
- [34] Marta Gomez-Barrero, Christian Rathgeb, Ulrich Scherhag, and Christoph Busch. Is your biometric system robust to morphing attacks? In *5th International Workshop on Biometrics and Forensics, IWBF 2017, Coventry, United Kingdom, April 4-5, 2017*, pages 1–6. IEEE, 2017. 2
- [35] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III*, volume 9907 of *Lecture Notes in Computer Science*, pages 87–102. Springer, 2016. 2
- [36] International Organization for Standardization. ISO/IEC DIS 30107-3:2016: Information Technology – Biometric presentation attack detection – P. 3: Testing and reporting, 2017. 7
- [37] ISO/IEC JTC1 SC37 Biometrics. ISO/IEC 29794-1:2016 Information technology - Biometric sample quality - Part 1: Framework. International Organization for Standardization, 2016. 4
- [38] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. 4, 7
- [39] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 4401–4410, 2019. 4
- [40] Vahid Kazemi and Josephine Sullivan. One millisecond face alignment with an ensemble of regression trees. In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, pages 1867–1874. IEEE Computer Society, 2014. 4
- [41] Ali Khodabakhsh, Ramachandra Raghavendra, Kiran B. Raja, Pankaj Shivdayal Wasnik, and Christoph Busch. Fake face detection methods: Can they be generalized? In *2018 International Conference of the Biometrics Special Interest Group, BIOSIG 2018, Darmstadt, Germany, September 26-28, 2018*, volume P-282 of *LNI*, pages 1–6. GI / IEEE, 2018. 5
- [42] Davis E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10:1755–1758, 2009. 4
- [43] Yehuda Lindell and Benny Pinkas. Secure multiparty computation for privacy-preserving data mining. *Journal of Privacy and Confidentiality*, 1(1), Apr. 2009. 3
- [44] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 3730–3738. IEEE Computer Society, 2015. 6
- [45] Ashwin Machanavajjhala, Daniel Kifer, John Abowd, Johannes Gehrke, and Lars Vilhuber. Privacy: Theory meets practice on the map. In *2008 IEEE 24th International Conference on Data Engineering*, pages 277–286, 2008. 3
- [46] Satya Mallick. Face morph using opencv — c++ / python. *LearnOpenCV*, 1(1), 2016. 4, 6, 7
- [47] Fabio Valerio Massoli, Fabio Carrara, Giuseppe Amato, and Fabrizio Falchi. Detection of face recognition adversarial attacks. *Comput. Vis. Image Underst.*, 202:103103, 2021. 1
- [48] Blaz Meden, Peter Rot, Philipp Terhörst, Naser Damer, Arjan Kuijper, Walter J. Scheirer, Arun Ross, Peter Peer, and Vitomir Struc. Privacy-enhancing face biometrics: A comprehensive survey. *IEEE Trans. Inf. Forensics Secur.*, 16:4147–4183, 2021. 3
- [49] Tom Neubert, Andrey Makrushin, Mario Hildebrandt, Christian Kraetzer, and Jana Dittmann. Extended stirtrace benchmarking of biometric and forensic qualities of morphed face images. *IET Biometrics*, 7(4):325–332, 2018. 7
- [50] Mei Ngan, Patrick Grother, Kayee Hanaoka, and Jason Kuo. Face Recognition Vendor Test (FRVT) Part 4: MORPH -

- Performance of Automated Face Morph Detection. *National Institute of Standards and Technology (NIST)*, 2021. [1](#)
- [51] European Parliament. European parliamentary research service, briefing public expectations and eu policies: Identifying the gaps, 7 2016. [3](#)
- [52] Alyssa Quek. *Facemorpher*. 2019. [7](#)
- [53] Ramachandra Raghavendra, Kiran B. Raja, and Christoph Busch. Detecting morphed face images. In *8th IEEE International Conference on Biometrics Theory, Applications and Systems, BTAS 2016, Niagara Falls, NY, USA, September 6-9, 2016*, pages 1–7. IEEE, 2016. [1](#), [2](#)
- [54] Ramachandra Raghavendra, Kiran B. Raja, Sushma Venkatesh, and Christoph Busch. Face morphing versus face averaging: Vulnerability and detection. In *2017 IEEE International Joint Conference on Biometrics, IJCB 2017, Denver, CO, USA, October 1-4, 2017*, pages 555–563. IEEE, 2017. [1](#), [2](#), [6](#)
- [55] Ramachandra Raghavendra, Kiran B. Raja, Sushma Venkatesh, and Christoph Busch. Transferable deep-cnn features for detecting digital and print-scanned morphed face images. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 1822–1830. IEEE Computer Society, 2017. [1](#), [2](#)
- [56] Raghavendra Ramachandra, Sushma Venkatesh, Kiran B. Raja, and Christoph Busch. Detecting face morphing attacks with collaborative representation of steerable features. In *Proceedings of 3rd International Conference on Computer Vision and Image Processing - CVIP 2018, Jabalpur, India, September 29 - October 1, 2018, Volume 1*, volume 1022 of *Advances in Intelligent Systems and Computing*, pages 255–265. Springer, 2018. [2](#), [5](#)
- [57] Eklavya Sarkar, Pavel Korshunov, Laurent Colbois, and Sébastien Marcel. Vulnerability analysis of face morphing attacks from landmarks and generative adversarial networks. *arXiv preprint*, Oct. 2020. [2](#), [4](#), [5](#), [7](#)
- [58] Ulrich Scherhag, Luca Debiase, Christian Rathgeb, Christoph Busch, and Andreas Uhl. Detection of face morphing attacks based on PRNU analysis. *IEEE Trans. Biom. Behav. Identity Sci.*, 1(4):302–317, 2019. [2](#)
- [59] Ulrich Scherhag, Ramachandra Raghavendra, Kiran B. Raja, Marta Gomez-Barrero, Christian Rathgeb, and Christoph Busch. On the vulnerability of face recognition systems towards morphed face attacks. In *5th International Workshop on Biometrics and Forensics, IWBF 2017, Coventry, United Kingdom, April 4-5, 2017*, pages 1–6. IEEE, 2017. [1](#), [2](#)
- [60] Ulrich Scherhag, Christian Rathgeb, Johannes Merkle, and Christoph Busch. Deep face representations for differential morphing attack detection. *IEEE Trans. Inf. Forensics Secur.*, 15:3625–3639, 2020. [2](#), [4](#)
- [61] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 2818–2826. IEEE Computer Society, 2016. [5](#)
- [62] The European Parliament and the Council of the European Union. Art. 17 of the general data protection regulation, 2016. [3](#)
- [63] The European Parliament and the Council of the European Union. Art. 7(3) of the general data protection regulation, 2016. [3](#)
- [64] The European Parliament and the Council of the European Union. Article 22(4) of the general data protection regulation, 2016. [2](#)
- [65] The European Parliament and the Council of the European Union. Article 26 of the general data protection regulation, 2016. [3](#)
- [66] The European Parliament and the Council of the European Union. Article 27(2) of the general data protection regulation, 2016. [2](#)
- [67] The European Parliament and the Council of the European Union. Article 30 of the general data protection regulation, 2016. [2](#)
- [68] The European Parliament and the Council of the European Union. Article 35(3) of the general data protection regulation, 2016. [2](#)
- [69] The European Parliament and the Council of the European Union. Article 37(1) of the general data protection regulation, 2016. [2](#)
- [70] The European Parliament and the Council of the European Union. Article 5(1)(a) of the general data protection regulation, 2016. [3](#)
- [71] The European Parliament and the Council of the European Union. Article 6(1)(a) of the general data protection regulation, 2016. [3](#)
- [72] The European Parliament and the Council of the European Union. Article 6(1)(b),(c),(d), and (f) of the general data protection regulation, 2016. [3](#)
- [73] The European Parliament and the Council of the European Union. Article 6(1)(e) of the general data protection regulation, 2016. [3](#)
- [74] The European Parliament and the Council of the European Union. Article 6(4) of the general data protection regulation, 2016. [2](#), [3](#)
- [75] The European Parliament and the Council of the European Union. Article 6(4)(c) of the general data protection regulation, 2016. [3](#)
- [76] The European Parliament and the Council of the European Union. Article 9 of the general data protection regulation, 2016. [2](#)
- [77] The European Parliament and the Council of the European Union. Article 9(2) of the general data protection regulation, 2016. [2](#)
- [78] The European Parliament and the Council of the European Union. Article 9(2)(a) of the general data protection regulation, 2016. [3](#)
- [79] The European Parliament and the Council of the European Union. Article 9(2)(b)-(i) of the general data protection regulation, 2016. [3](#)
- [80] The European Parliament and the Council of the European Union. Article 9(2)(j) of the general data protection regulation, 2016. [3](#)

- [81] The European Parliament and the Council of the European Union. Article 9(4) of the general data protection regulation, 2016. [2](#)
- [82] The European Parliament and the Council of the European Union. 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 (General Data Protection Regulation), 2016. [2](#)
- [83] Sushma Venkatesh, Haoyu Zhang, Raghavendra Ramachandra, Kiran B. Raja, Naser Damer, and Christoph Busch. Can GAN generated morphs threaten face recognition systems equally as landmark based morphs? - vulnerability and detection. In *8th International Workshop on Biometrics and Forensics, IWBIF 2020, Porto, Portugal, April 29-30, 2020*, pages 1–6. IEEE, 2020. [2](#), [7](#)
- [84] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z. Li. Learning face representation from scratch. *CoRR*, abs/1411.7923, 2014. [2](#)
- [85] Haoyu Zhang, Sushma Venkatesh, Raghavendra Ramachandra, Kiran Bylappa Raja, Naser Damer, and Christoph Busch. MIPGAN - generating strong and high quality morphing attacks using identity prior driven GAN. *IEEE Trans. Biom. Behav. Identity Sci.*, 3(3):365–383, 2021. [2](#)
- [86] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Process. Lett.*, 23(10):1499–1503, 2016. [5](#)
- [87] Yuanhan Zhang, Zhenfei Yin, Yidong Li, Guojun Yin, Junjie Yan, Jing Shao, and Ziwei Liu. Celeba-spoof: Large-scale face anti-spoofing dataset with rich annotations. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XII*, volume 12357 of *Lecture Notes in Computer Science*, pages 70–85. Springer, 2020. [2](#)