

## Worst-Case Morphs: a Theoretical and a Practical Approach

Una M. Kelly,<sup>1</sup> Luuk Spreeuwers<sup>1</sup>, Raymond Veldhuis<sup>1</sup>

**Abstract:** Face Recognition (FR) systems have been shown to be vulnerable to morphing attacks. We examine exactly how challenging morphs can become. By showing a worst-case construction in the embedding space of an FR system and using a mapping from embedding space back to image space we generate images that show that this theoretical upper bound can be approximated if the FR system is known. The resulting morphs can also successfully fool unseen FR systems and are useful for exploring and understanding the weaknesses of FR systems. Our method contributes to gaining more insight into the vulnerability of FR systems.

**Keywords:** Biometrics, Morphing Attack Detection, Face Recognition, Vulnerability of Biometric Systems.

### 1 Introduction

Several publications have shown that Face Recognition (FR) systems and humans are vulnerable to morphing attacks. A morph is an image that contains sufficient identity information from the faces of two different individuals for it to be accepted as a match when it is compared to an image of either contributing identity. This can lead to security risks in e.g. border control, since a criminal could travel using the identity document of an accomplice. Several methods to address such attacks have been proposed. In most cases, the morphs on which Morphing Attack Detection (MAD) methods are trained and tested are generated in-house, and are usually created by detecting corresponding landmarks in the faces, warping the images to an average geometry and then blending the pixel values [FFM14, RKB17, MW18].

The fact that MAD methods are often tested on datasets with similar characteristics as the sets used for training them may lead to overfitting to certain dataset-specific characteristics - especially if only one morphing algorithm was used [SRB18, SSV18]. In many countries someone applying for an identity document can provide their own printed passport photo, or even upload a photo digitally (e.g. Ireland, Estonia). This means that before a passport photo is stored in an electronic Machine Readable Travel Document (eMRTD), it may have been printed and scanned or digitally manipulated.

There are several other factors that affect how morphed images differ, leading to variation in the performance of MAD methods and vulnerability of FR systems. These include the

---

<sup>1</sup>Data Management and Biometrics, University of Twente, Enschede, The Netherlands, {u.m.kelly, l.j.spreeuwers, r.n.j.veldhuis}@utwente.nl

morphing algorithm used, which landmark detector was used, or whether landmarks were selected manually, printing and scanning, the pairs of images selected for morphing, etc. There may be other tools for generating morphs of which the research community is not yet aware. Therefore, it is necessary to explore other potential methods and better understand the weaknesses of FR systems. Benchmarks for validating MAD methods - such as [Ra20, NI, BO] - can be extended and improved by evaluating morphs generated with different methods. We contribute to this active field of research by providing a new morphing method, and by showing that there is an upper limit on how difficult morphs can become.

We investigate exactly how challenging morphs can in theory and in practice become for FR systems. Given an FR system and two images of two different people, the *worst-case* morph is the image that is most similar to both images, according to this FR system. We train a neural network to generate images that approximate such worst-case morphs and examine how well our approximations of worst-case morphs can fool an FR system, and whether this extends to other systems. Our results show it is possible to approximate the worst case when the (weights of the) FR system are known, which means that morphing attacks can become even more challenging than landmark- or GAN-based morphs. In black box scenarios, when the network weights are not accessible, our approximations of the worst case are less challenging, but still pose a significant threat.

Our proposed system differs from GAN-based morphs [Zh21, Ve20a] since GANs are trained to generate images from an embedding space (different from the embedding space of an FR system) that look like real images, while we directly work with the embedding space of an FR system and its associated (dis)similarity measure to approximate the most challenging morph, and we train without the adversarial constraint of GANs.

Our main contributions are: a theoretical framework that when given an FR system and two images can be used to define a *worst-case* embedding, a practical demonstration of finding images that (approximately) correspond to this worst-case embedding, showing that these can be significantly more challenging than landmark- or GAN-based morphs, examining the vulnerability of other FR systems, and providing a new method for generating morphs that leads to more variation in morph datasets, which may help researchers more accurately validate MAD methods in future work and uncover vulnerabilities of FR systems.

## 2 Background

Research on variation in morphing algorithms includes: post-processing landmark-based morphs to mask effects caused by the morphing process [SHE19], a model to simulate the effects of printing and scanning [FFM21] and considering the influence of ageing on morphing attacks [Ve20b]. GANs were used in [Ve20a, Zh21, Da21] in an attempt to create a different type of morph, which was shown to be able to fool FR systems, if not as consistently as landmark-based morphs. Since GANs are difficult to train more stable methods would be useful [Cr18]. The lack of variation in morphing techniques is addressed in [Sc20]. A method for MAD is presented and evaluated on morphs created using different algorithms, which are all landmark-based. Printed-and-scanned morphs are included, but GAN morphs or other methods are not taken into consideration.




Fig. 1: Overview of our morphing approach. Here,  $x_1, x'_1$  are images of one person and  $x_2, x'_2$  are images of a different person. While  $x_{\text{morph}}$  visually looks like  $x_1$ , according to the FR system it is a successful match with both  $x'_1$  and  $x'_2$ . Dashed lines represent comparisons by the FR system.

### 3 Proposed System

Traditionally, landmark-based morphing attempts to create images in image space that are similar to two given images. Instead, we use the embedding space of an FR system, since this should contain structured information on the similarity of the identity information present in images. Let  $f$  be the function that describes an FR system's mapping from the image space  $X$  to the embedding space  $Z$ , i.e.  $f : X \rightarrow Z$  (Fig. 1). Let  $d$  be the dissimilarity score that calculated the dissimilarity of embedding pairs in  $Z$ . We define the *worst-case embedding* for two images  $x_1$  and  $x_2$  as

$$z^* := \operatorname{argmin}_{z \in Z} (\max [d(z, f(x_1)), d(z, f(x_2))]). \quad (1)$$

For example, if  $S$  returns cosine similarity, then  $S(z_1, z_2) = \cos(\theta)$ , where  $\theta$  is the angle between  $z_1$  and  $z_2$ , see Fig. 2. In that case  $z^*$  is any  $z$  for which  $S(z_1, z) = S(z, z_2) = \cos(\theta/2)$ . If an FR system uses similarity scores, defined by a function  $S$ , then  $S$  replaces  $d$ , argmin is replaced by argmax, and max is replaced by min.




Fig. 2: The worst-case embedding  $z^*$  when  $d$  denotes euclidean distance (left) and angle (right). An image that maps to  $z^*$  is even more challenging than a landmark- ( $z_{\text{lm}}$ ) or GAN-based morph ( $z_{\text{GAN}}$ ).

A worst-case morph is an image  $x^*$  for which  $f(x^*) = z^*$ . We approximate  $x^*$  using a decoder  $D$  that maps from  $Z$  to  $X$ , reversing the mapping of the FR system. Since  $f$  is a many-to-one mapping there may be many different approximative inverse mappings for one given embedding. We choose to train  $D$  to reconstruct  $x_1$  when given  $z^*$  as input, i.e.  $D$  learns to approximate the mapping  $g^* : Z \rightarrow X$ , where  $g^*(z^*) = x_1$  and  $f(g^*(z^*)) = z^*$ .  $D$  generates images that to humans look like  $x_1$ , i.e. like *one* of the input identities (e.g. the accomplice), but to the FR system are a morph of  $x_1$  and  $x_2$ . According to the FR system the resulting image contains identity features of both accomplice and criminal.

Worst-case morphs

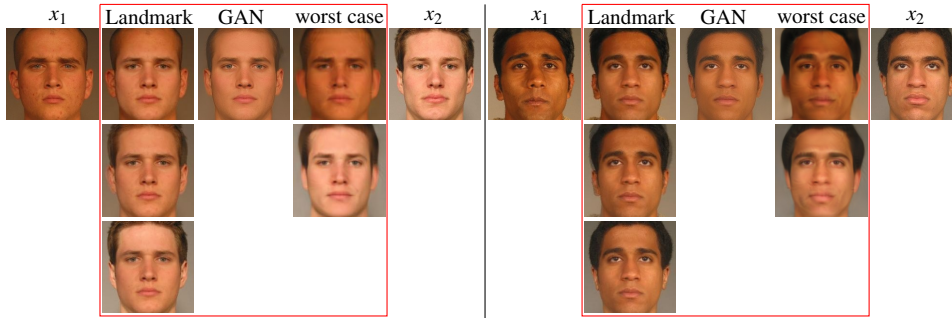


Fig. 3: Landmark-based morphs (spliced into  $x_1$ , full morph, spliced into  $x_2$ ), GAN-based morph and worst-case approximations, with  $z_{\text{Enc}}$  (top) from  $x_1$  and with  $z_{\text{Enc}}$  from  $x_2$  (bottom).

Instead of training a decoder to generate morphed images, we could also use optimisation on an FR system. In that case, the FR latent embeddings corresponding to an image are manipulated by calculating gradients and computing adversarial perturbations. The disadvantage of learning (i.e. training a decoder or other network) over optimisation (i.e. gradient descent) is that the decoder is trained to generate many different images, while optimisation applies an adversarial perturbation tailored specifically to one image. The advantage of learning over optimisation is that there is no need for retraining or calculating image-specific gradients every time a new image is manipulated (e.g. from a new dataset).

### 3.1 Morph generation

For our experiments we use MobileFaceNet, which on LFW achieves an accuracy of 99.55% [Ch18]. For images  $x_1$  and  $x_2$ , this FR system returns 128-dimensional embeddings  $z_1 = f(x_1)$  and  $z_2 = f(x_2)$ , where  $f$  describes the mapping defined by the neural network from the image space to the embedding space:

$$X = [0, 1]^{n_c \times w \times h} \quad \text{and} \quad Z = \{z \in \mathbb{R}^{128} \mid \|z\| = 1\},$$

where  $n_c = 3$  and  $w \times h = 112 \times 112$ . It can be used as a face verification system by calculating the angle  $\theta$  between  $z_1$  and  $z_2$ .  $x_1$  and  $x_2$  are not accepted as a match if  $\theta > t$ , where  $t$  is the decision threshold. In this case, the worst-case embedding  $z^*$  that minimises Eq. 1 is the embedding that lies on the unit hypersphere exactly between  $z_1$  and  $z_2$ , i.e. a worst-case morph is any image  $x^*$  for which  $f(M^*) = z^* = \frac{z_1 + z_2}{\|z_1 + z_2\|}$ . Scaling is necessary to ensure that  $z^*$  lies on the unit hypersphere.

An embedding  $z = f(x)$  should contain information essential to the identity of  $x$ , but some information such as background or expression is not relevant to identity and should therefore not be contained in  $z$ . This means that the FR latent embedding alone is not sufficient to reconstruct  $x$ , so we train a supporting encoder network that also compresses image information, including other attributes. We train the decoder network using both worst-case embeddings of the FR system and embeddings of the encoder as input.

---

**Algorithm 1** Training procedure.  $\gamma_1 = 1, \gamma_2 = 0.1$ 


---

 $\theta_D, \theta_{\text{Enc}} \leftarrow$  initialize network parameters

**repeat**
 $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}$  and  $\mathbf{z}^{(i)} = f(\mathbf{x}^{(i)})$ ,  $i = 1, \dots, N$   $\triangleright$  Draw  $N$  samples from the dataset and get FR emb.

 $\mathbf{z}^{*(i)} = \frac{\mathbf{z}^{(i)} + \mathbf{z}^{(j)}}{\|\mathbf{z}^{(i)} + \mathbf{z}^{(j)}\|}$ ,  $i = 1, \dots, N$ ,  $j = 2, \dots, N, 1$   $\triangleright$  Translate batch to get pairs for worst-case emb.

 $\mathbf{z}_{\text{Enc}}^{(i)} = f_{\text{Enc}}(\mathbf{x}^{(i)})$ ,  $i = 1, \dots, N$   $\triangleright$  Get Encoder emb.

 $\mathbf{x}_{\text{morph}}^{(i)} = D(\mathbf{z}^{(i)}, \mathbf{z}_{\text{Enc}}^{(i)})$ ,  $i = 1, \dots, N$   $\triangleright$  Generate morphs

 $\mathbf{z}_{\text{morph}}^{(i)} = f(\mathbf{x}_{\text{morph}}^{(i)})$ ,  $i = 1, \dots, N$   $\triangleright$  Get FR embeddings of morphs.

 $\mathcal{L} = \gamma_1 \frac{1}{N} \sum_{i=0}^N \text{MSE}(\mathbf{x}^{(i)}, \mathbf{x}_{\text{morph}}^{(i)}) + \gamma_2 \frac{1}{N} \sum_{i=0}^N d(\mathbf{z}^{*(i)}, \mathbf{z}_{\text{morph}}^{(i)})$   $\triangleright$  Compute loss

 $\theta_D \leftarrow \theta_D - \nabla_{\theta_D} \mathcal{L}$   $\triangleright$  Gradient update on decoder

 $\theta_{\text{Enc}} \leftarrow \theta_{\text{Enc}} - \nabla_{\theta_{\text{Enc}}} \mathcal{L}$   $\triangleright$  Gradient update on encoder

**until** convergence

---

Let  $f_{\text{Enc}}$  be the mapping of the supporting encoder, i.e.  $f_{\text{Enc}}(x) = z_{\text{Enc}}$ . Furthermore, let  $g$  denote the decoder mapping back to image space. Let  $d$  be the dissimilarity score function used by the FR system to compare latent embeddings, which here is the angle between embeddings. We train  $D$  to approximate  $g^*$  by minimising the following loss

$$\begin{aligned} \mathcal{L} &= \gamma_1 \mathcal{L}_{\text{pixel}} + \gamma_2 \mathcal{L}_{\text{latent}} \\ &= \gamma_1 \mathbb{E}_x [\|x_{\text{morph}} - x_1\|_2^2] + \gamma_2 \mathbb{E}_x [d(z_{\text{morph}}, z^*)], \end{aligned} \quad (2)$$

where  $\mathcal{L}_{\text{pixel}}$  ensures that the output images are visually similar to the input images, and  $\mathcal{L}_{\text{latent}}$  encourages the corresponding FR embeddings to be close to the worst case. The architecture and hyperparameters of this network can be found in the appendix. We train the network for 200 epochs, after which the losses no longer significantly decrease.

We use a dataset of in total 21,772 facial images from [Ph05] and separate them into 18,143 training and 3,629 validation images, with no overlap in identities. We create three sets of morphs using the validation set: approximations of worst-case morphs, landmark-based morphs and GAN-based morphs using MIPGAN-I [Zh21]. We select pairs of similar identities and faces with neutral expression, resulting in 506 pairs for morphing. For each pair  $(x_1, x_2)$  we create two worst-case approximations, three landmark morphs (one full morph and two spliced morphs), and one GAN morph, see Fig. 3.

We improve each morph generated with the trained Decoder by freezing the network’s weights and performing gradient descent on  $\mathcal{L}_{\text{latent}}$  (Eq. 4) for  $N$  iterations to improve input selection, where  $z_1^* = z^* = \frac{z_1 + z_2}{\|z_1 + z_2\|}$  and  $z_i^* = z_{i-1}^* - \nabla \mathcal{L}_{\text{latent}}(z_{\text{morph}})$ . The improved morph is obtained by forwarding  $(z_N^*, z_{\text{Enc}})$  through the decoder. We set  $N = 200$ , after which the white box FR scores still decrease, but not the scores for black box FR systems.

## 4 Results & Discussion

We compare the vulnerability of FR systems to landmark-based, GAN-based morphing attacks and worst-case approximations. We examine the morphs with the same deep-

## Worst-case morphs

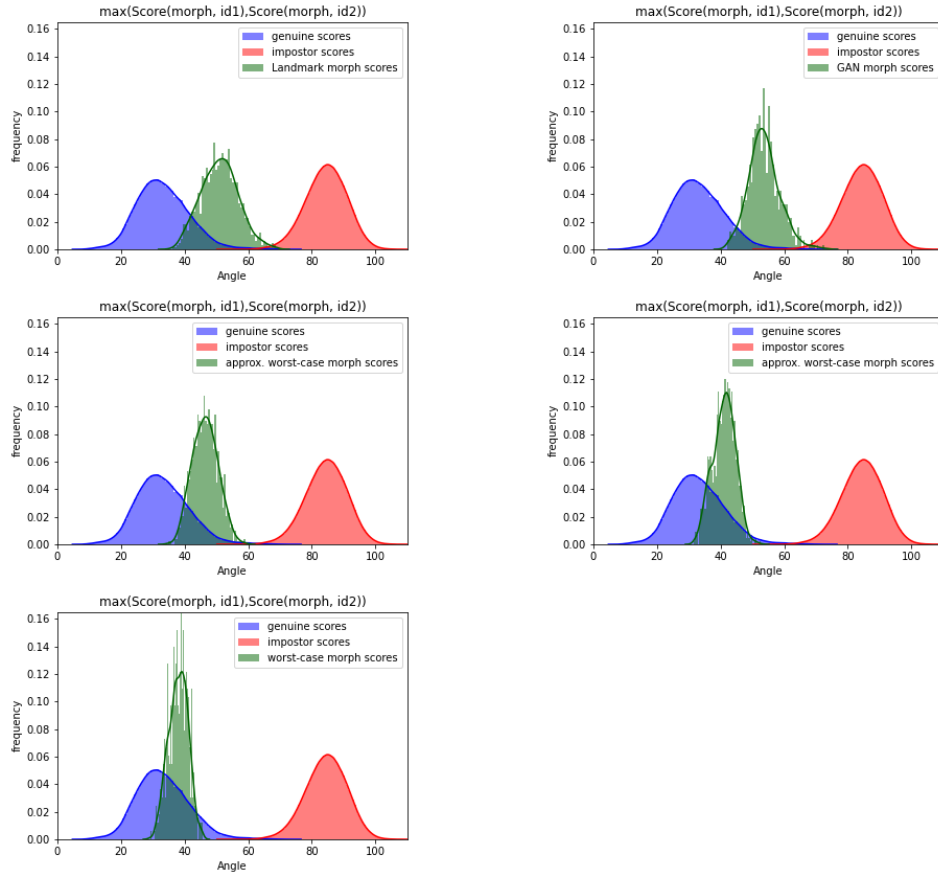


Fig. 4: Vulnerability of MobileFaceNet to: landmark morphs, MIPGAN-I morphs, approximations of worst-case morphs, improved approximations after optimisation (200 iterations per image), and worst-case embeddings (theoretical upper bound). Blue histograms describe genuine comparison scores, impostor comparison scores are red. For each morph we compute the angles for probe images of each contributing identity. The green histograms describe the largest of these two angles. Our morphs are extremely difficult to distinguish from genuines using FR comparison scores.

learning-based FR system used for training the inverse network and with two other FR systems, where we calculate the Mated Morph Presentation Match Rate ( $MMPMR(t)$ ), which is the proportion of (morphing) attacks for which both contributing identities are considered a match by the FR system when using a threshold  $t$  [Sc17]. For each FR system we set  $t$  such that the false non-match rate is minimal while the false match rate  $< 0.1\%$ .

The histograms of MobileFaceNet comparison scores are shown in Fig. 4. The  $MMPMR$  for the worst-case approximations is 100%, even higher than the 97.3% for landmark-based morphs (Tab. 1). This means that every worst-case approximation contains enough identity information of both contributing identities to fool the FR system.

Tab. 1: MMPMR (%).

	Landmark	GAN	Worst-case	Worst-case approx.	Improved worst-case approx.
<i>MobileFaceNet</i>	97.3	96.3	100.0	100.0	100.0
<i>VGG16</i> (black box)	63.0	65.0	100.0	33.6	43.1
<i>ArcFace</i> (black box)	85.4	66.3	100.0	47.9	61.1

In Fig. 4 we can see that the worst-case approximations are very close to the actual worst case, especially after optimising each morph separately. Table 1 shows that this does not necessarily hold for other FR systems. This is not entirely surprising, since our method is trained to generate approximations of worst-case morphs specifically for one FR system, while the worst-case morphs for a different FR system might be different. Nonetheless, a significant proportion of the morphs is still accepted as a match with both contributing identities. Most importantly, we have shown that it is possible to generate images that are close to the worst case, and therefore believe that researchers should not focus exclusively on landmark- and GAN-based morphs, since more challenging morphs may exist.

## 5 Conclusion & Future Work

We introduced the concept of worst-case morphs, which are morphs that are most similar to the contributing identities according to an FR system. We trained a decoder network to generate images that approximate such worst-case morphs. These were extremely successful at fooling the FR system that was used for training and significantly more challenging than landmark- or GAN-based morphs. We also showed that generating morphs that are close to the worst case for FR systems under black box settings is still an open problem. This might be improved by minimising the distance to the worst-case embedding for several FR systems simultaneously. Since the quality of morphs depends on the image quality of genuine images, pair selection, and other factors, a good measure to compare the quality of different morphs could be measuring the distance to the worst-case embedding.

The morphs we generate may be helpful in understanding the vulnerabilities of FR systems and may also offer some insight into the robustness of MAD techniques. Since the morphs generated using the decoder network are closer to the worst case than landmark- or GAN-based morphs, even before applying optimisation to improve each morph individually, challenging morphs can be generated in a computationally inexpensive manner. Furthermore, our approach has the advantage of being more stable to train than e.g. GANs. We also studied the potential of our method for use in black box settings.

## References

- [BO] Bologna Online Evaluation Platform (BOEP) - Morph Attack Detection Evaluation, <https://biolab.csr.unibo.it/fvcongoing/UI/Form/BOEP.aspx>, accessed: 05-10-2021.

- 
- [Ch18] Chen, Sheng; Liu, Yang; Gao, Xiang; Han, Zhen: MobileFaceNets: Efficient CNNs for Accurate Real-time Face Verification on Mobile Devices. CoRR, abs/1804.07573, 2018.
- [Cr18] Creswell, A. et al.: Generative Adversarial Networks: An Overview. IEEE Signal Processing Magazine, 35(1):53–65, 2018.
- [Da21] Damer, N. et al.: , ReGenMorph: Visibly Realistic GAN Generated Face Morphing Attacks by Attack Re-generation, 2021.
- [FFM14] Ferrara, M.; Franco, A.; Maltoni, D.: The magic passport. In: IEEE International Joint Conference on Biometrics. pp. 1–7, 2014.
- [FFM21] Ferrara, M.; Franco, A.; Maltoni, D.: Face morphing detection in the presence of printing/scanning and heterogeneous image sources. IET Biometrics, 10(3):290–303, 2021.
- [MW18] Makrushin, A.; Wolf, A.: An Overview of Recent Advances in Assessing and Mitigating the Face Morphing Attack. In: 2018 26th European Signal Processing Conference (EU-SIPCO). pp. 1017–1021, 2018.
- [NI] NIST FRVT MORPH, [https://pages.nist.gov/frvt/html/frvt\\_morph.html](https://pages.nist.gov/frvt/html/frvt_morph.html), accessed: 05-10-2021.
- [Ph05] Phillips, P. J. et al.: Overview of the face recognition grand challenge. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05). volume 1, pp. 947–954 vol. 1, 2005.
- [Ra20] Raja, K. et al.: Morphing Attack Detection – Database, Evaluation Platform and Benchmarking. arXiv, 2020.
- [RKB17] Robertson, D.; Kramer, R.; Burton, A.: Fraudulent ID using face morphs: Experiments on human and automatic recognition. PLOS ONE, 12:e0173319, 03 2017.
- [Sc17] Scherhag, U. et al.: Biometric Systems under Morphing Attacks: Assessment of Morphing Techniques and Vulnerability Reporting. In: 2017 International Conference of the Biometrics Special Interest Group (BIOSIG). pp. 1–7, 2017.
- [Sc20] Scherhag, U.; Kunze, J.; Rathgeb, C.; Busch, C.: Face morph detection for unknown morphing algorithms and image sources: a multi-scale block local binary pattern fusion approach. IET Biometrics, 9(6):278–289, 2020.
- [SHE19] Seibold, C.; Hilsmann, A.; Eisert, P.: Style Your Face Morph and Improve Your Face Morphing Attack Detector. In: 2019 International Conference of the Biometrics Special Interest Group (BIOSIG). pp. 1–6, 2019.
- [SRB18] Scherhag, U.; Rathgeb, C.; Busch, C.: Performance variation of morphed face image detection algorithms across different datasets. In: 2018 International Workshop on Biometrics and Forensics (IWBF). pp. 1–6, 2018.
- [SSV18] Spreeuwens, L.; Schils, M.; Veldhuis, R.: Towards Robust Evaluation of Face Morphing Detection. In: 2018 26th European Signal Processing Conference (EUSIPCO). pp. 1027–1031, 2018.
- [Ve20a] Venkatesh, S. et al.: 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). pp. 1–6, 2020.
- [Ve20b] Venkatesh, S.; Raja, K.; Ramachandra, R.; Busch, C.: On the Influence of Ageing on Face Morph Attacks: Vulnerability and Detection. In: 2020 IEEE International Joint Conference on Biometrics (IJB). pp. 1–10, 2020.



- [Zh21] Zhang, H.; Venkatesh, S.; Ramachandra, R.; Raja, K.; Damer, N.; Busch, C.: MIPGAN - Generating Strong and High Quality Morphing Attacks Using Identity Prior Driven GAN. IEEE Transactions on Biometrics, Behavior, and Identity Science, 3(3):365–383, 2021.

## Hyperparameters

Operation	Kernel	Stride	BN	Bias	Nonlinearity	Output Size
$f_{\text{Enc}}(\mathbf{x})$					LReLU(0.02)	
5× Convolution	$5 \times 5$	$2 \times 2$	✓	False	LReLU(0.02)	$7 \times 7 \times 64$
$g(\mathbf{z})$						
Fully connected	-	-	✗	True	LReLU(0.02)	$1 \times 1 \times 512$
Fully connected	-	-	✗	True	LReLU(0.02)	$7 \times 7 \times 448$
Concat. with $f_{\text{Enc}}(\mathbf{x})$			✓	False	LReLU(0.02)	$7 \times 7 \times 512$
4× Upsample & Conv.	$3 \times 3$	$1 \times 1$	✓	False	LReLU(0.02)	$112 \times 112 \times 32$
Conv.	$3 \times 3$	$1 \times 1$	✓	False	LReLU(0.02)	$112 \times 112 \times 32$
Upsample & Conv.	$3 \times 3$	$1 \times 1$	✓	Untied	Sigmoid	$224 \times 224 \times 3$
Optimizer	Adam( $\alpha = 10^{-4}$ , $\beta_1 = 0.0$ , $\beta_2 = 0.9$ ), batch size = 64					
Weight init.	Isotropic gaussian ( $\mu = 0$ , $\sigma = 0.01$ )					
Untied bias init.	Sigmoid $^{-1}(\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i)$ , the average of $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ .					