UNIVERSIDADE DE COIMBRA

Impact of Image Context for Deep Learning Face Morphing Attack Detection

Joana Alves Pimenta

Supervision of Prof. Nuno Gonçalves, MsC Iurii Medvedev

July, 2023

Outline

01

Context and Motivation

02

Goals and Contributions

03

Methodology

04

Experiments and Results

Outline

01

Context and Motivation

02

Goals and Contributions

03

Methodology

04

Experiments and Results

Currently, FRSs are used in a variety of applications, such as document security, **border control systems (ABC gates)**

At the same time **human face** can undergo several **modifications**

Increase the vulnerability of FRSs

Exposing them to **fraud attempts and criminal attacks** (presentation attacks).

Face Morphing

Combine the facial features of two (or more) images to create a synthetic image that **incorporates characteristics from both** faces

Face Morphing

The resulting image can be **easily confused** with the faces of two or more individuals, as it incorporates a combination of their facial characteristics.

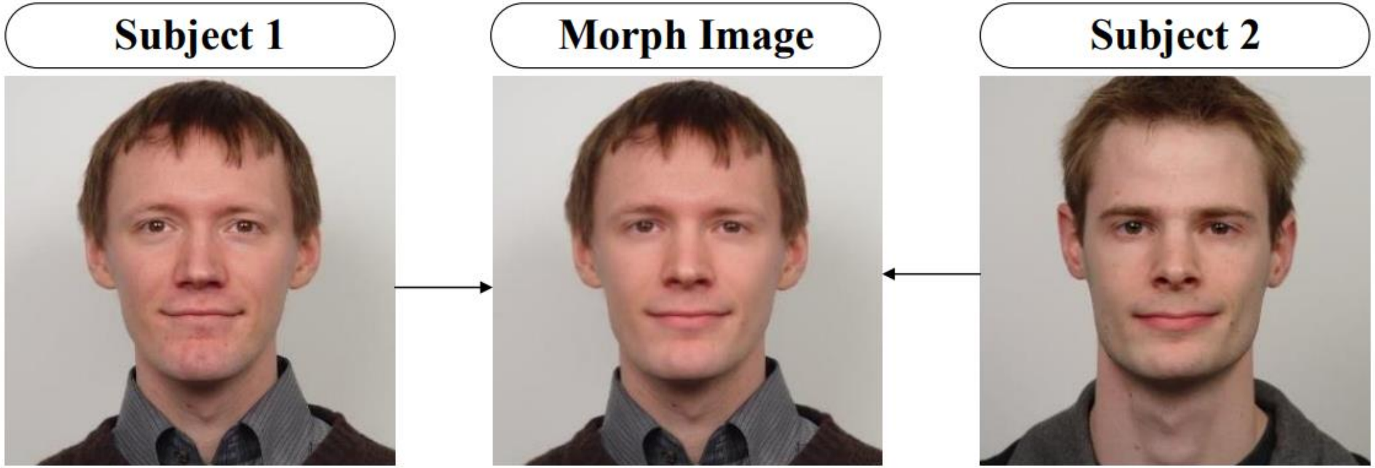

Schematic representation of face morphing between two subjects face images. Images from IMM dataset

Face Morphing

Allows to obtain an image of an individual **that appears to be real but, in fact, does not exist**

Which one of these images is real?

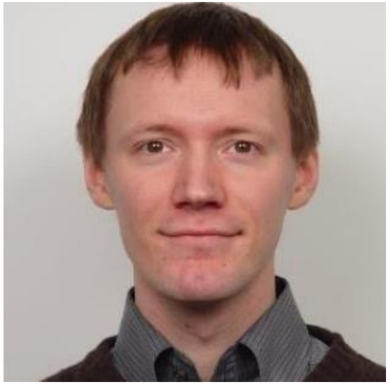

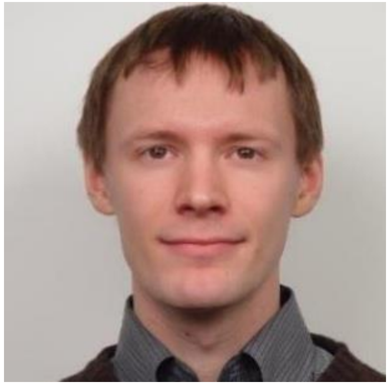

Face Morphing

Makes it possible to obtain a **legitimate ID using false information**

Passport application process

One individual can **impersonate another**, thus violating the principle of **exclusive ownership**

Same document may be validated for two or more persons

Albanian individual who attempted to pass through **LISBON** border control **using a Slovenian passport**

Accomplice

Morphed image presented on the document

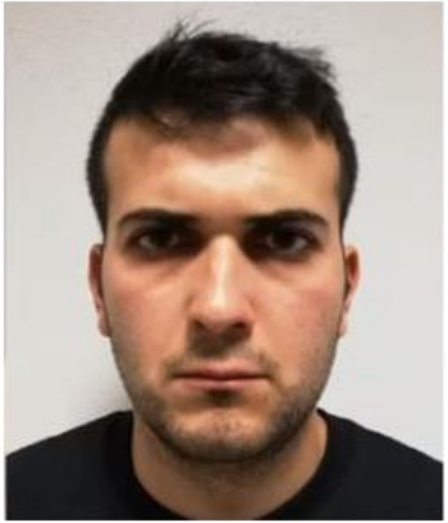

Person attempting to pass

Real case example of a face morphing attack

~~Face Morphing~~

Development of new **protection solutions**

Morphing Attack Detection (MAD) techniques

In this dissertation

- Contribute to advances in the study of MAD by **investigating whether or not image context influences the detection.**

Motivation

- MAD performance is **influenced** by factors like **alignment** and pre-processing techniques.
- **Face alignment impact contextual information captured in input images** which can affect the detection algorithm's performance.

Outline

01

Context and Motivation

02

Goals and Contributions

03

Methodology

04

Experiments and Results

Goals

Exploring how different face image **alignment settings can impact** the amount of context captured in the input image.

Finding the best context properties for detection, i.e., **defining optimal alignment settings** for face morphing detection.

Contributions

- Creation of a **large dataset** that adheres to the **ICAO** standards through the combination and pre-processing of multiple datasets.

- **Generation of a morphed dataset** using both landmark-based and StyleGAN based approaches.

- Investigation of the relationship between image context and MAD to **identify the most effective context properties for detection.**

- Formulation and implementing several strategies for MAD.

- **Submission of a paper** titled ”Impact of Image Context for Single Deep Learning Face Morphing Detection” in the **BIOSIG 2023 conference**

Outline

01

Context and Motivation

02

Goals and Contributions

03

Methodology

04

Experiments and Results

General Pipeline

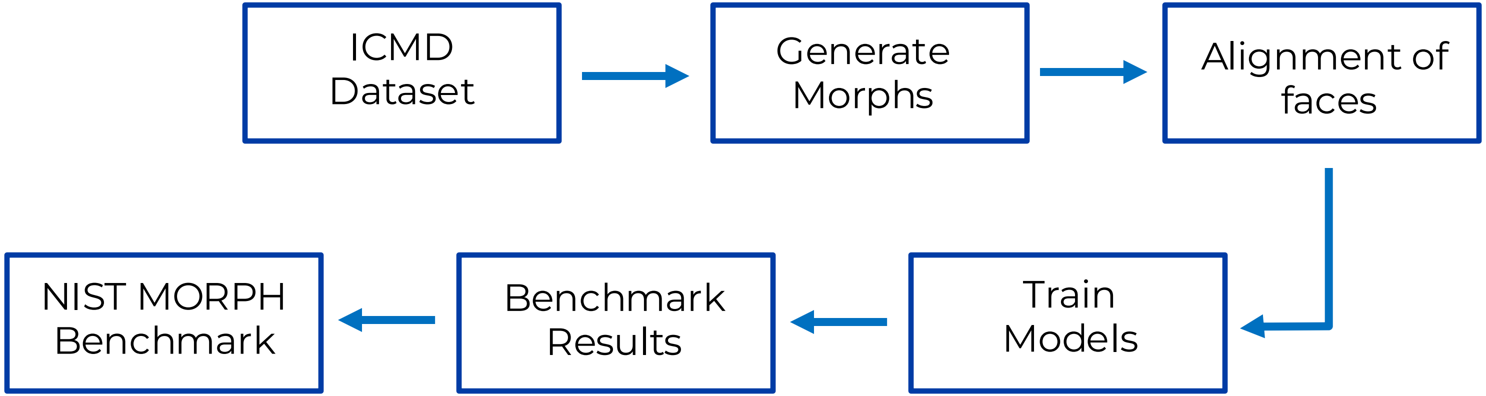

Source Data Curating

Lack of large *ICAO compliant* datasets (public ID document compliant)

Aggregation of several datasets →

ICMD dataset
~50k images of
~2500 individuals

Datasets Names	Images Number	Identities Number
FRGC	4007	466
XM2VTS	1180	295
ND Twins	24 050	435
FERET	11 000	994
AR	4000	126
PICS	17 122	141
FEI	2800	200
IMMFFD	1440	12
GTBD	150	50

Summary table of datasets used

There may be identical images between datasets

Morphed Image Generation

Landmark-based
Approach

StyleGAN-based
Approach

Selfmorphs
assumed as bonafide

Image-level

Representation-level

To focus morphing detection on the behavior of **deep face features**, **rather** than to **detecting artefacts**

Landmark-based Approach

Face landmark alignment → Image Warping → Blending → Context Restoring

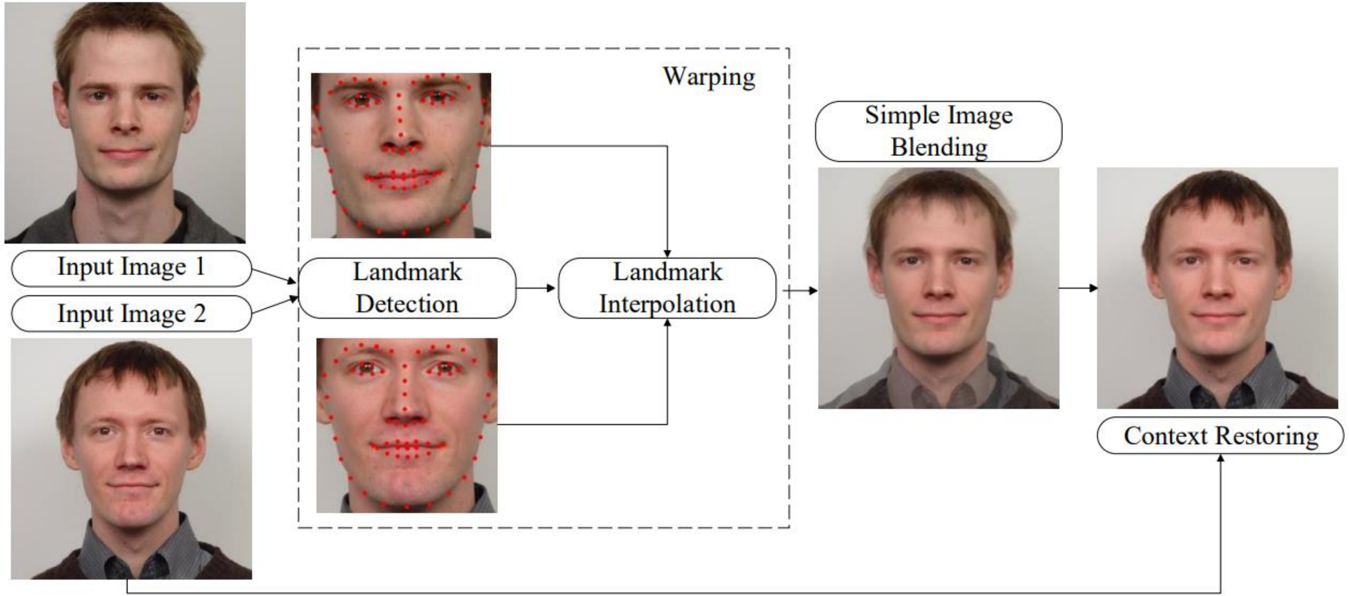

Landmark face morphing generation pipeline

Triangulation is performed not only with the reference points but also **using the borders** of the image.

Distortion problems when the face is not centrally located in the image

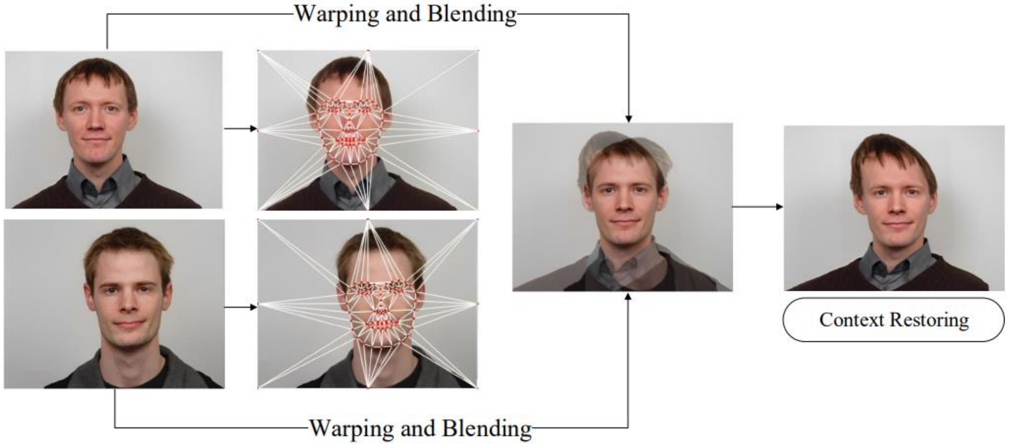

Distortion problem in landmark-based approach

Cropping the original images

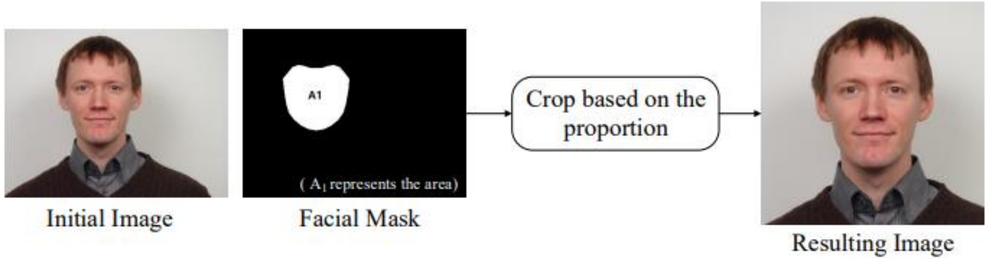

Pre-processing pipeline to deal with the distortion problem

StyleGAN-based Approach

Latent Embeddings → Linear Interpolation → Generator → Context Restoring

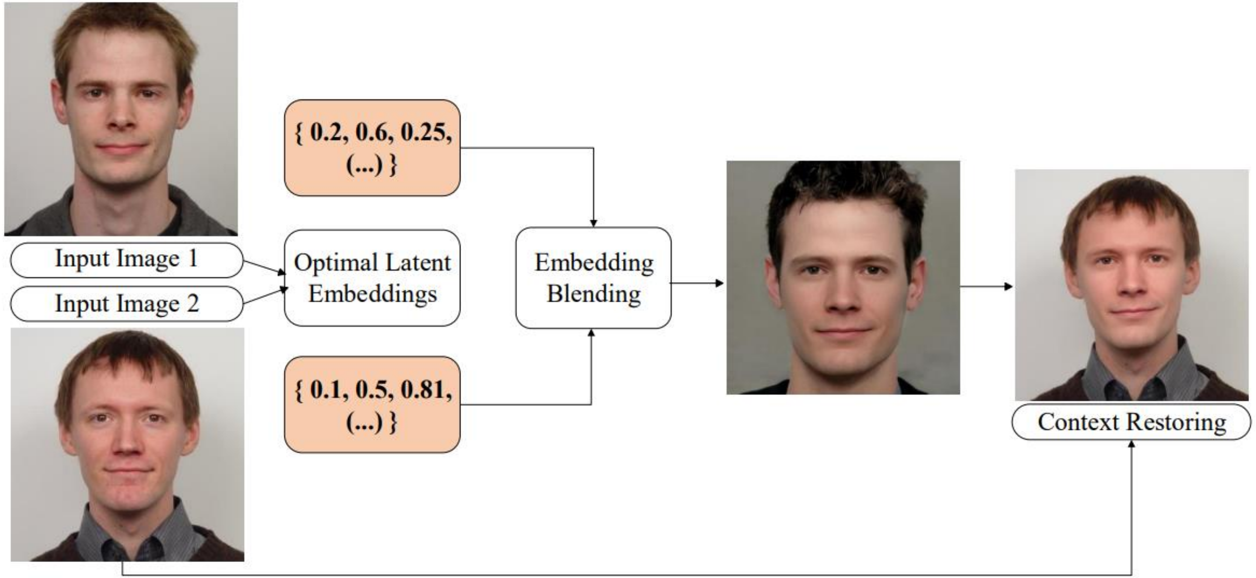

StyleGAN interpolation and morph generation pipeline

Morph Paring Approach

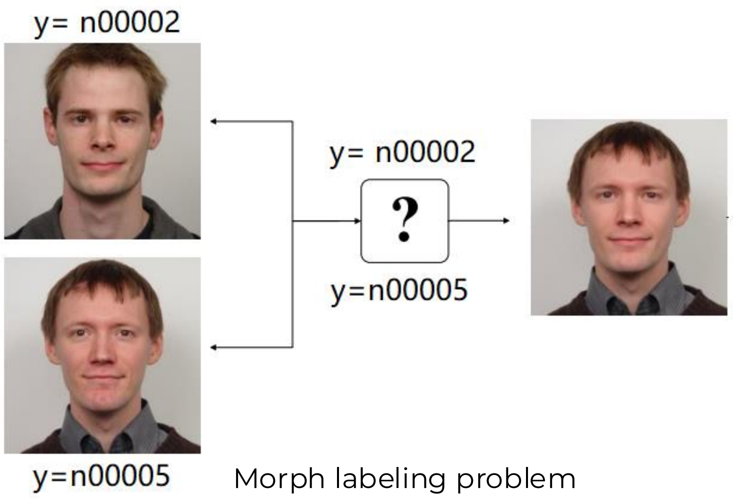

Morph labeling problem

Our approach considers
face morphing detection
from the perspective of face recognition
MorDeephy

Ambiguity and classification confusion.

The resulting morphed image belongs to both source identities, i.e., it simultaneously **has two class labels associated with it.**

Separate the total list of identities **into two halves**

Pairs of images were generated from each half

Each generated image was **labeled based on the corresponding sub-list for further classification.**

Assure that morphed **combinations of a particular identity** are classified **similarly**

Medvedev, I.; Shadmand, F.; Gonçalves, N.: MorDeephy: Face Morphing Detection via Fused Classification. In: Proceedings of the 12th ICPRAM. SciTePress, pp. 193–204, 2023.

Alignment Settings

Different alignment conditions were defined in **order to vary the relationship between the face and the background of the image**.

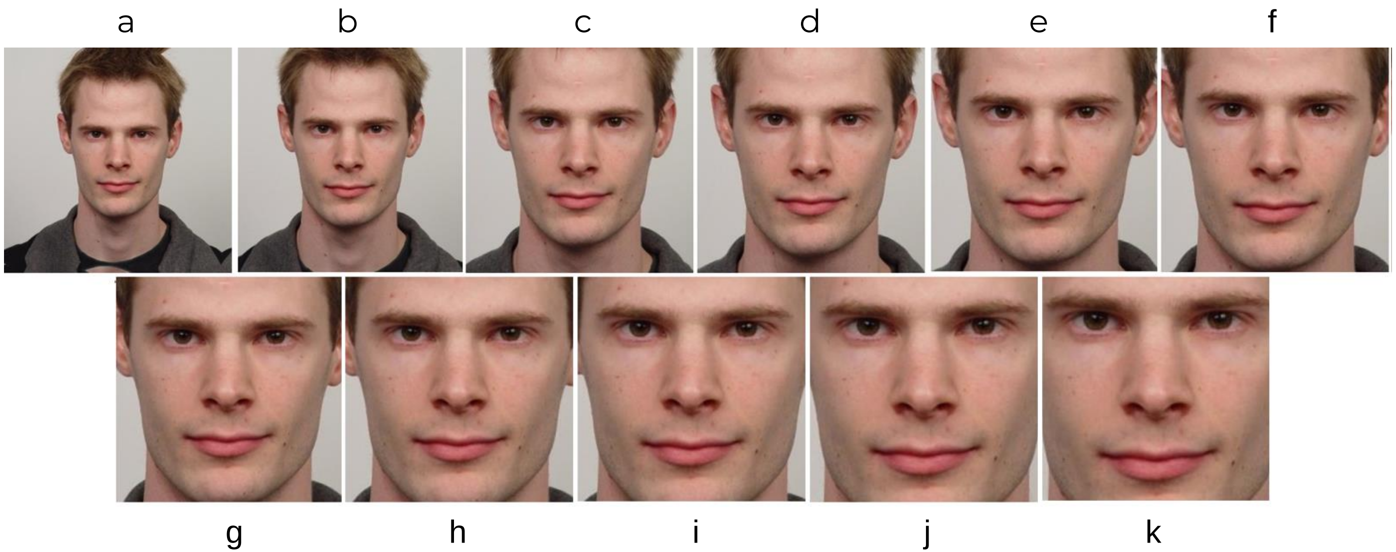

Face image aligned according to the different alignment settings

5 Facial landmarks
({left eye}, {right eye}, {nose}, {left mouth corner}, {right mouth corner})

→

Rigid transformation
minimize the coordinate distance between those **five facial landmarks and a predefined target set**

→

Scaling this target coordinate set using several scale factors

Face Morphing Attack Detection

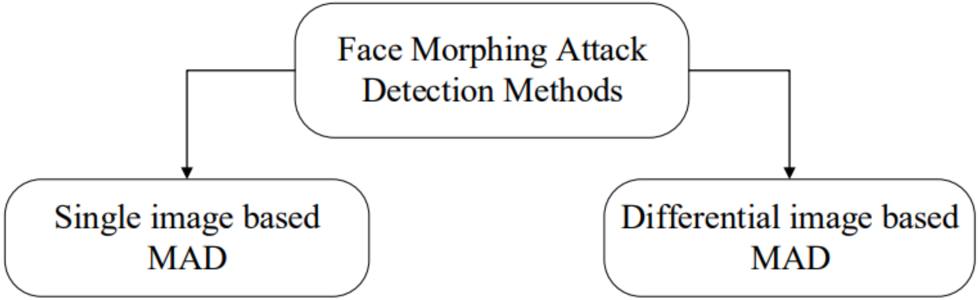

The S-MAD case uses **a single image**

Non-reference method

Real-life scenarios: Initial passport application

The D-MAD case involves **comparing** the **test image** to a **reference image**

***Reference-based method*.**

Real-life scenarios: ABC gates

Single Image MAD Approach

Fused Classification Approach

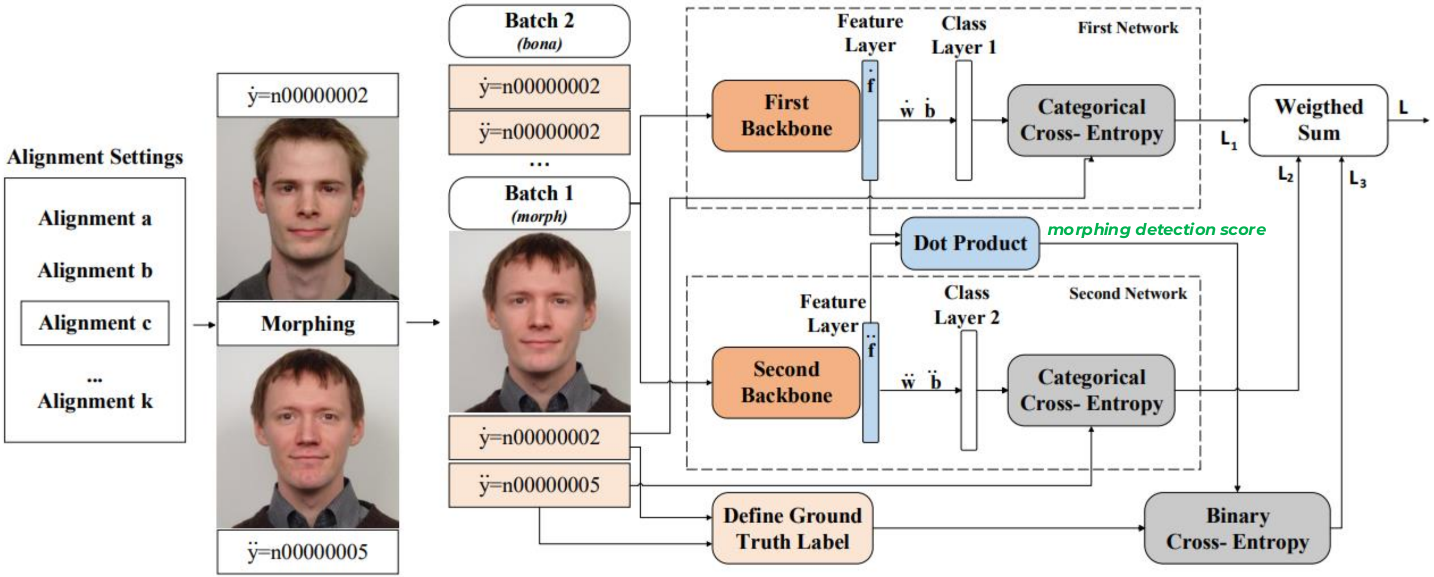

S-MAD model schema for fused classification approach.

Consider ***bonafide*** samples **similarly** and **morphed** images **differently**

Learn **high-level features** by performing classification tasks

The identity parts act as a **regularization**

→

Retains the facial **discriminability** of feature layers (unique and distinguishing facial characteristics)

The features are then compared with a **similarity measure (dot product)**

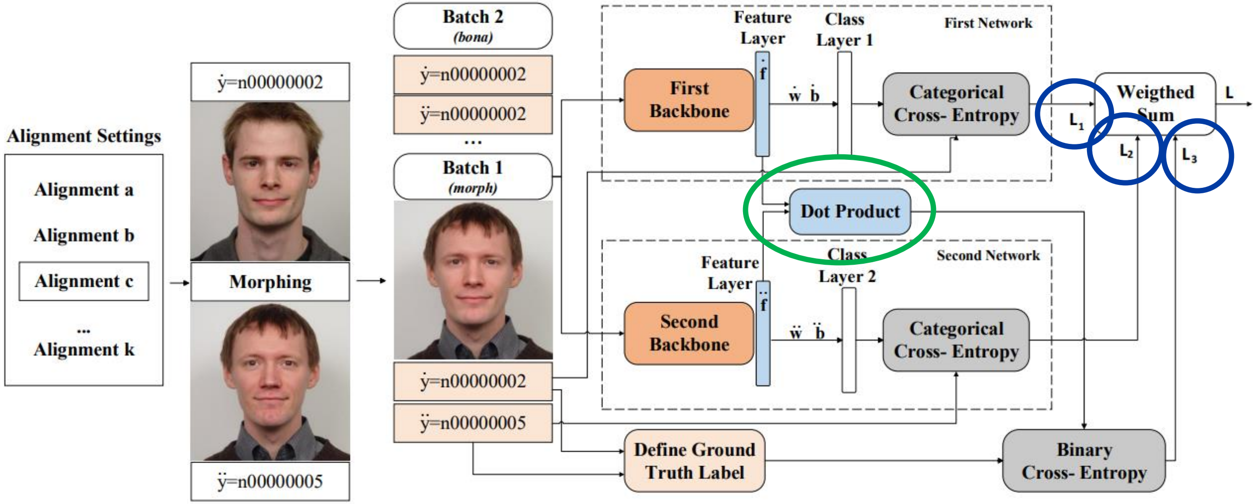

S-MAD model schema for fused classification approach.

$$L = \alpha_1 L_1 + \alpha_2 L_2 + \beta L_3$$

By **minimizing** this loss configuration, the model **learns discriminative facial features for effective morphing detection**.

Softmax loss Functions

$$L_1 = -\frac{1}{N}\sum_i^N \log\left(\frac{e^{\dot{W}^T_{\dot{y}_i}\dot{f}_i + \dot{b}_{\dot{y}_i}}}{\sum_j^C e^{\dot{f}_{\dot{y}_j}}}\right)$$

$$L_2 = -\frac{1}{N}\sum_i^N \log\left(\frac{e^{\ddot{W}^T_{\ddot{y}_i}\ddot{f}_i + \ddot{b}_{\ddot{y}_i}}}{\sum_j^C e^{\ddot{f}_{\ddot{y}_j}}}\right)$$

Dot Product + Sigmoid Function → Binary cross-entropy

$$L_3 = -\frac{1}{N}\sum_i^N t\log\frac{1}{1+e^{-\dot{f}\cdot\ddot{f}}} + (1-t)\log\left(1-\frac{1}{1+e^{-\dot{f}\cdot\ddot{f}}}\right)$$

$$t = 1 - |sgn(\dot{y}_i - \ddot{y}_i)|$$

1 – bonafides (equal labels)
0 – morphs (different labels)

Ground- Truth label

Binary Classification Approach

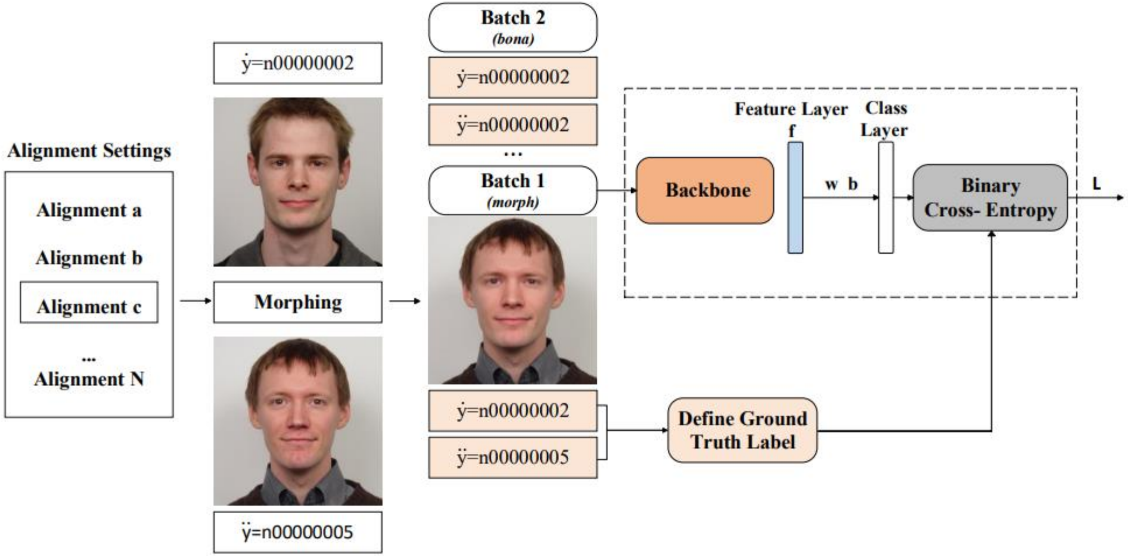

S-MAD model schema for binary classification approach.

This modification involves **removing the identity classification** component present in the fused classification approach.

Differential Image MAD Approach

Fused Classification Approach

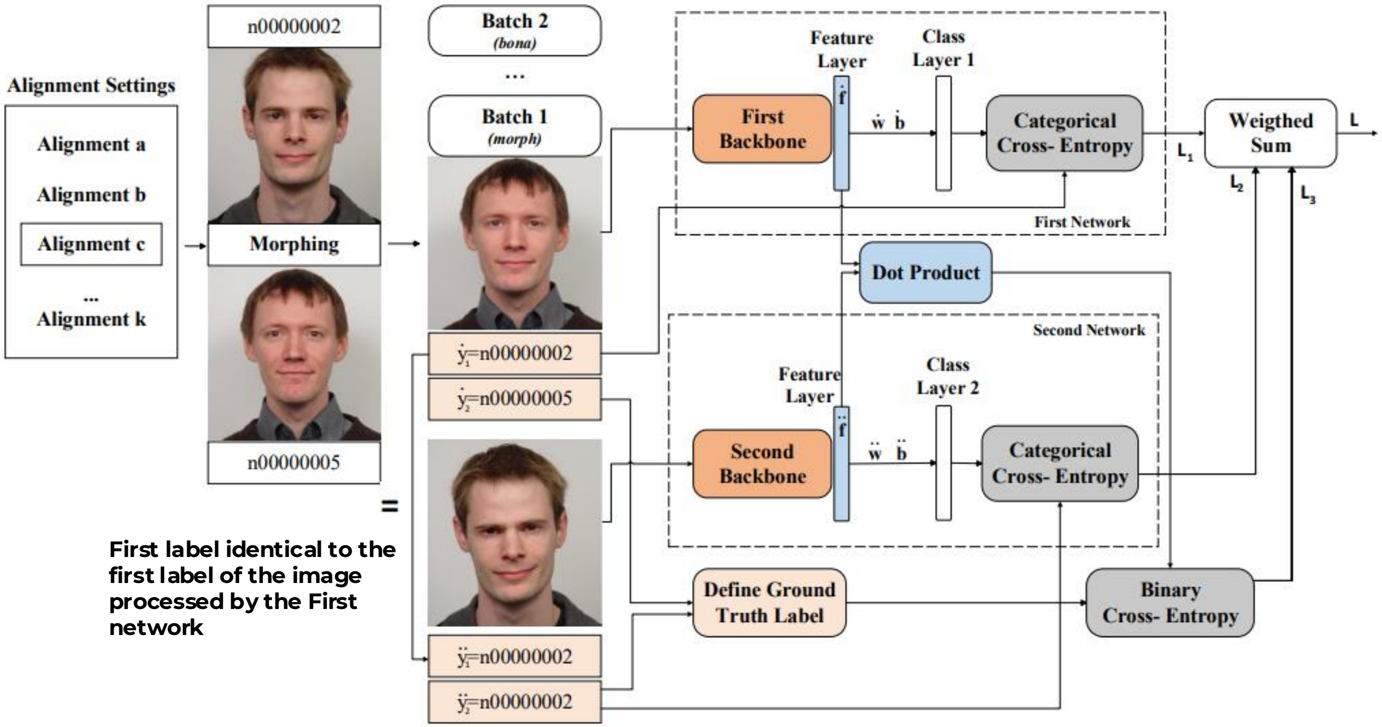

First label identical to the first label of the image processed by the First network

D-MAD model schema for fused classification approach.

First and Second networks receive **two images** instead of the same single image

First network receives the ”test image”(morph or not)

Second network receives the ”live capture image” (always bonafide).

Grad-CAM Approach

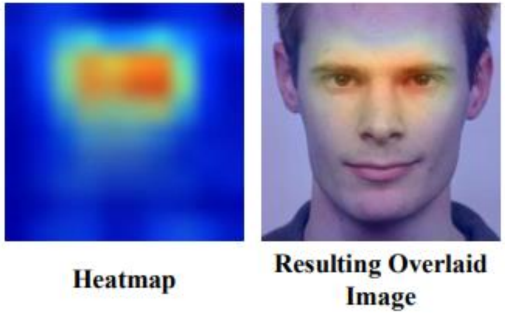

Grad-CAM sample heatmap and its overlaid sample image.

Obtain a final heatmap that **highlights** the important regions of the input image, providing valuable **insights into the decision-making process of the model**

Gradient for **real binary classification** (morph and non-morph)

Single network backbone

Obtained **directly** by tracking the activations of a certain **class score**

Two network backbones

Obtained by taking the gradient of **dot product between the feature embeddings** with respect to the activations of the last convolutional layers of both networks

The maps **were also divided according to morph or *bonafide*** *cases*

Average maps by protocol and respective alignment condition

Average of the Gradient Intensity Ratio (AGIR) *foreground/background*

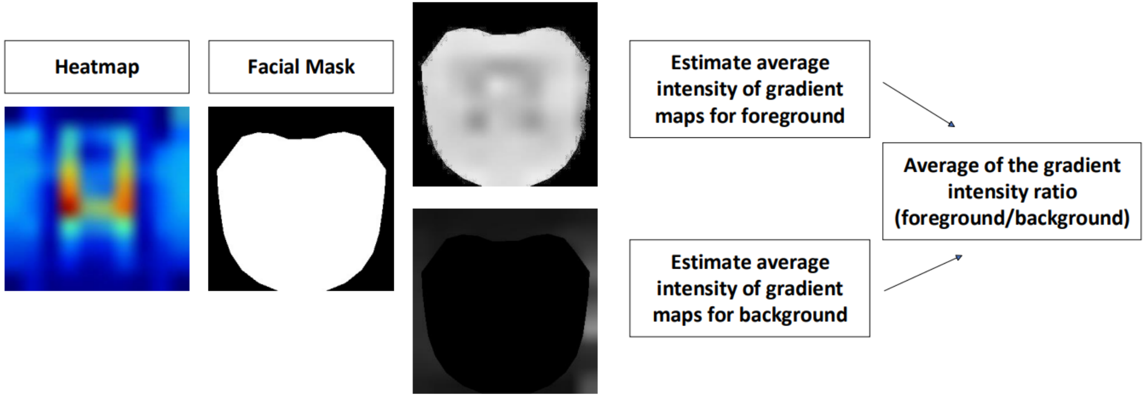

Schematic representation of the methodology to obtain the average intensity of gradient maps for the foreground and background and the respective ratio.

The average gradient intensity for each region was computed **using only the non-zero pixels**

Ratio between the two average intensities.

Outline

Training Settings

Backbone architecture

EfficientNetB3, initialized with **weights pre-trained on the ImageNet** returning in the end 512 deep features

The batch size was set to 28 images.
The optimizer employed was **stochastic gradient descent (SGD)** with a momentum parameter of 0.9.

Hyperparameters

- Choose the **appropriate balance** among the components of the **loss function (α_1, α_2, and β)**

$$L = \alpha_1 L_1 + \alpha_2 L_2 + \beta L_3$$

- Epochs Number
- Initial Learning Rate

Benchmarking

Insights about the performance of a particular model.

Public Benchmarks

FVC-onGoing MAD

FRVT NIST MORPH

Submission restrictions →

Open-source morphing benchmarking utilities (MorDeephy)

Benchmarking Protocols

asml

- ~ **2k morphs (S-MAD)**
- ~ **4.3k morphs (D-MAD)**

opencv

- ~ **1.3k morphs (S-MAD)**
- ~ **2.4k morphs (D-MAD)**

facemorpher

- ~ **2k morphs (S-MAD)**
- ~ **2.4k morphs (D-MAD)**

webmorph

- ~ **1k morphs (S-MAD)**
- ~ **2.4k morphs (D-MAD)**

stylegan

- ~ **2k morphs (S-MAD)**
- ~ **2.4k morphs (D-MAD)**

real

- ~ **3k morphs (S-MAD)**

APCER

Attack Presentation Classification Error Rate
Proportion of morph images that are wrongly classified as *bonafide* (user insecurity).

Morph Miss Rate

BPCER

Bonafide Presentation Classification Error Rate
Proportion of *bonafide* images that are wrongly classified as morphs (user inconvenience).

False Detection Rate

These two metrics are typically presented as: **BPCER@APCER**

↓

Goal
Balance between APCER (Security) and BPCER (Convenience)
Lowest possible values

Benchmark Results

S-MAD Binary Classification Model

Alignments	BPCER@APCER=δ											
	Protocol-asml		Protocol-facemorpher		Protocol-opencv		Protocol-stylegan		Protocol-webmorph		Protocol-real	
	δ =0.1	δ =0.01	δ =0.1	δ =0.01	δ =0.1	δ =0.01	δ =0.1	δ =0.01	δ =0.1	δ =0.01	δ =0.1	δ =0.01
a	0.199	0.622	0.125	0.558	0.199	0.663	0.663	0.663	0.523	0.663	0.191	0.568
b	0.143	0.380	0.131	**0.387**	0.144	0.440	0.586	0.586	0.340	0.586	**0.144**	**0.396**
c	0.365	0.630	0.331	0.675	0.320	0.676	0.676	0.676	0.489	0.676	0.351	0.630
d	0.236	0.511	0.161	0.549	0.161	0.489	0.623	0.623	0.436	0.623	0.246	0.511
e	**0.141**	**0.348**	**0.102**	0.532	**0.080**	**0.424**	0.710	0.710	**0.321**	0.641	0.194	0.463
f	0.199	0.455	0.127	0.551	0.125	0.533	0.675	0.675	0.328	0.579	0.215	0.478
g	0.158	0.373	0.106	0.532	0.209	0.532	0.586	0.586	0.348	0.586	0.175	0.411
h	0.330	0.580	0.138	0.682	0.093	0.486	0.724	0.724	0.486	0.724	0.306	0.577
i	0.214	0.408	0.174	0.476	0.149	0.442	**0.573**	**0.573**	0.396	**0.573**	0.212	0.430
j	0.221	0.465	0.187	0.596	0.141	0.457	0.776	0.776	0.475	0.682	0.233	0.504
k	0.243	0.498	0.194	0.557	0.146	0.513	0.794	0.794	0.467	0.707	0.262	0.573

BPCER@APCER=δ	a	b	c	d	**e**	**f**	**g**	h	i	j	k
δ=0.1	0.317	0.248	0.442	0.320	0.258	0.278	0.263	0.346	0.286	0.338	0.351
δ=0.01	0.623	0.463	0.660	0.551	0.520	0.545	0.503	0.629	0.484	0.580	0.607

Optimal range : the alignment settings between ***e*** and ***g***
Potential optimal case : alignment setting ***e***

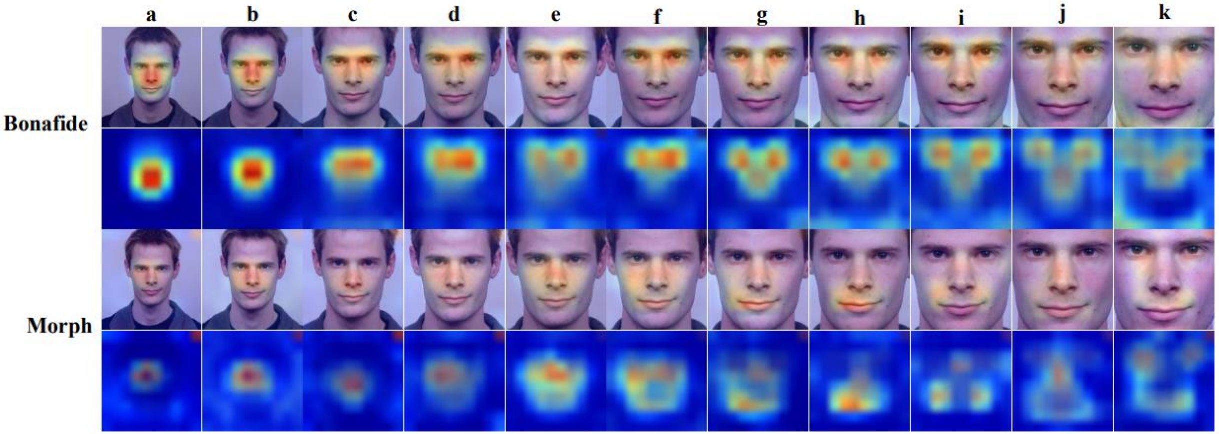

Grad-CAM heatmaps across all the alignment settings

Alignments	AGIR morph values						AGIR value for bona	Average AGIR value for morphs
	asml	facemorpher	opencv	stylegan	webmorph	real		
a	1.754	2.824	2.394	0.546	0.776	1.841	4.481	1.689
b	1.844	3.094	3.298	0.687	1.107	2.114	3.354	2.024
c	1.797	3.352	3.790	0.879	1.574	2.129	2.017	2.254
d	2.164	3.610	3.725	1.140	1.238	2.469	1.811	2.391
e	3.586	4.458	4.796	1.443	2.060	3.587	1.589	3.322
f	2.775	3.675	4.393	1.835	1.902	2.957	1.397	2.922
g	2.067	3.726	3.526	1.419	1.246	2.459	1.774	2.407
h	1.765	3.393	4.405	2.079	1.376	2.079	1.534	2.516
i	2.337	3.717	4.281	1.461	1.770	2.579	1.221	2.691
j	2.475	3.550	3.868	1.369	1.581	2.532	1.238	2.563
k	1.629	2.067	2.586	1.270	1.313	1.857	0.951	1.787

Summary table for the AGIR values in the different protocols, as well as the average value for the morphs.

The **face/foreground** is mostly dominantly activated across all the alignment settings

Benchmark Results

S-MAD Fused Classification Model

Alignments	BPCER@APCER=δ											
	Protocol-asml		Protocol-facemorpher		Protocol-opencv		Protocol-stylegan		Protocol-webmorph		Protocol-real	
	δ=0.1	δ=0.01	δ=0.1	δ=0.01	δ=0.1	δ=0.01	δ=0.1	δ=0.01	δ=0.1	δ=0.01	δ=0.1	δ=0.01
a	0.159	0.689	0.187	0.517	0.239	0.599	0.842	0.946	0.606	0.885	0.137	0.608
b	0.063	0.495	0.072	0.646	0.099	0.658	0.671	0.946	0.702	0.964	0.081	0.427
c	0.125	0.467	0.215	0.588	0.240	0.566	0.694	0.884	0.541	0.859	0.167	0.455
d	0.040	0.374	0.102	0.558	0.103	0.568	0.574	0.835	0.305	0.781	0.113	0.421
e	0.162	0.580	0.149	0.582	0.177	0.602	0.566	0.767	0.605	0.870	0.138	0.549
f	0.184	0.530	0.180	0.488	0.175	0.451	0.582	0.788	0.517	0.785	0.158	0.479
g	**0.034**	**0.233**	**0.025**	0.701	**0.037**	0.701	0.487	0.875	**0.216**	0.788	**0.072**	**0.322**
h	0.168	0.642	0.168	0.535	0.165	0.599	0.536	0.850	0.542	0.854	0.138	0.594
i	0.046	0.255	0.036	**0.365**	0.044	**0.390**	**0.305**	**0.583**	0.246	**0.554**	0.094	0.365
j	0.287	0.630	0.268	0.585	0.262	0.564	0.844	0.959	0.697	0.907	0.228	0.574
k	0.193	0.652	0.253	0.745	0.262	0.792	0.825	0.953	0.674	0.915	0.178	0.611

BPCER@APCER=δ	a	b	c	**d**	**e**	**f**	**g**	**h**	**i**	j	k
δ=0.1	0.361	0.281	0.330	0.206	0.299	0.299	0.145	0.286	0.129	0.431	0.398
δ=0.01	0.732	0.690	0.636	0.589	0.658	0.586	0.603	0.679	0.418	0.703	0.778

Optimal range : the alignment settings between ***d*** and ***i***

Potential optimal case : alignment setting ***g***

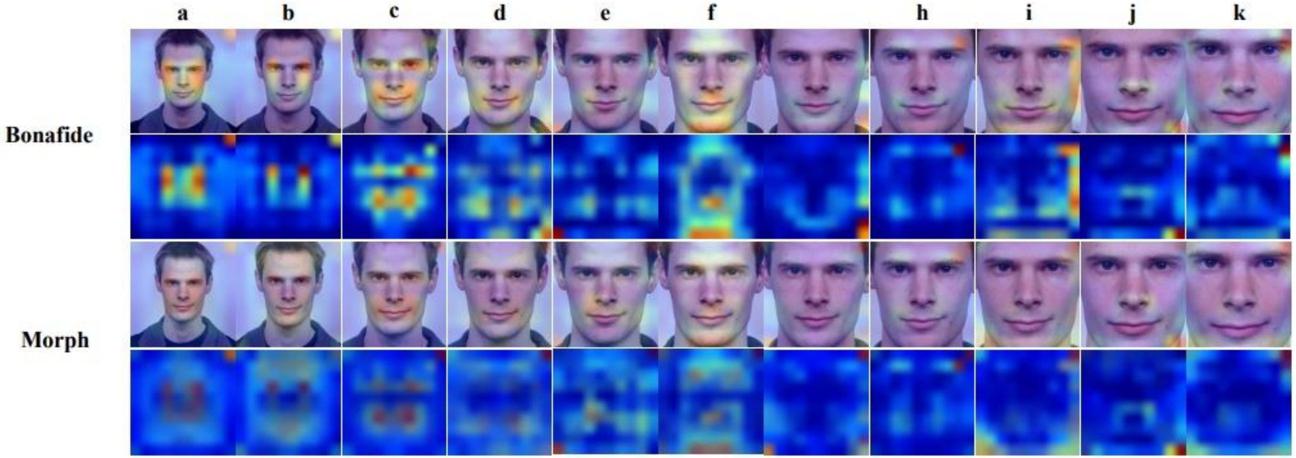

Grad-CAM heatmaps across all the alignment settings

Alignments	AGIR morph values						AGIR value for bona	Average AGIR value for morphs
	asml	facemorpher	opencv	stylegan	webmorph	real		
a	1.125	1.120	1.123	1.985	2.083	1.209	2.348	1.441
b	0.970	0.861	0.892	2.354	2.077	1.088	1.934	1.374
c	1.403	1.516	1.545	2.687	2.555	1.565	2.671	1.879
d	0.853	0.786	0.798	1.644	1.219	0.954	1.385	1.042
e	1.001	0.920	0.943	1.413	1.394	0.992	1.055	1.111
f	1.184	1.200	1.243	2.064	1.816	1.245	1.836	1.458
g	0.679	0.625	0.617	1.012	1.000	0.729	1.094	0.777
h	1.208	1.127	1.181	1.782	1.772	1.175	1.509	1.374
i	0.640	0.580	0.599	1.267	0.906	0.692	1.331	0.781
j	0.910	0.843	0.906	1.203	1.246	0.962	1.433	1.011
k	0.500	0.516	0.542	0.934	0.700	0.549	0.824	0.624

Summary table for the AGIR values in the different protocols, as well as the average value for the morphs.

The detection focuses **mainly on the face region** and, in many cases, on the **intersection regions** between foreground and background

Benchmark Results

D-MAD Fused Classification Model

Alignments	BPCER@APCER=δ									
	Protocol-asml		Protocol-facemorpher		Protocol-opencv		Protocol-stylegan		Protocol-webmorph	
	δ =0.1	δ =0.01	δ =0.1	δ =0.01	δ =0.1	δ =0.01	δ =0.1	δ =0.01	δ =0.1	δ =0.01
a	0.013	0.194	0.013	0.250	0.019	0.206	0.506	0.725	0.244	0.806
b	**0.000**	0.394	0.006	0.356	0.013	0.400	0.563	0.888	0.288	0.713
c	0.063	0.344	0.138	0.644	0.181	0.644	0.656	0.919	0.306	0.831
d	0.100	0.613	0.144	0.613	0.175	0.588	0.625	0.894	0.381	0.769
e	0.138	0.688	0.219	0.744	0.194	0.744	0.544	0.781	0.544	0.831
f	0.013	0.263	**0.000**	**0.206**	**0.006**	**0.206**	**0.188**	**0.525**	0.244	**0.569**
g	0.056	0.363	0.050	0.494	0.069	0.550	0.381	0.706	0.319	0.644
h	0.056	0.475	0.031	0.531	0.056	0.544	0.419	0.713	0.300	0.719
i	0.150	0.531	0.088	0.531	0.069	0.494	0.381	0.738	0.575	0.900
j	0.013	**0.144**	0.013	0.506	0.013	0.438	0.319	0.794	**0.163**	0.738
k	0.044	0.288	0.044	0.369	0.038	0.313	0.500	0.868	0.288	0.856

BPCER@APCER=δ	a	b	c	d	e	f	g	h	i	j	k
δ=0.1	0.159	0.174	0.269	0.285	0.327	0.09	0.175	0.172	0.252	0.103	0.182
δ=0.01	0.4362	0.550	0.676	0.695	0.757	0.350	0.551	0.596	0.639	0.524	0.538

Optimal range : the alignment settings between ***f*** and ***h***
Potential optimal case : alignment setting ***f***

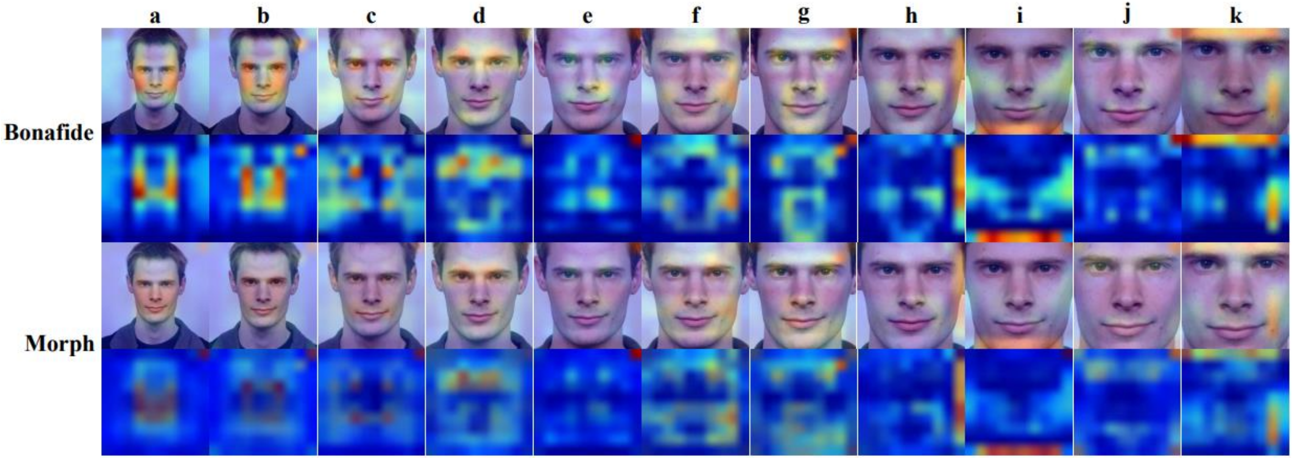

Grad-CAM heatmaps across all the alignment settings

Alignments	*AGIR* morph value					AGIR value for bona	Average AGIR value for morphs
	asml	facemorpher	opencv	stylegan	webmorph		
a	1.604	1.331	1.446	2.064	2.719	1.865	1.833
b	1.061	1.070	1.085	2.321	2.012	2.306	1.509
c	0.975	0.950	0.952	1.127	1.179	1.127	1.037
d	1.192	1.148	1.123	1.828	1.807	1.811	1.420
e	1.275	1.135	1.149	2.183	1.967	2.231	1.542
f	1.018	0.965	1.029	1.504	1.274	1.080	1.158
g	1.056	0.975	1.967	1.808	1.575	1.967	1.476
h	0.789	0.747	0.791	0.863	1.030	0.865	0.844
i	1.249	1.148	1.110	1.405	1.398	1.423	1.262
j	0.733	0.664	0.656	1.315	1.179	1.196	0.909
k	0.842	0.867	0.896	0.699	0.746	0.688	0.810

Summary table for the AGIR values in the different protocols, as well as the average value for the morphs.

Similar to S-MAD fused classification, the detection **focuses primarily on the face region** and, in many instances, on the intersections of the foreground and background

Discussion on the Results

- There is **possibly a region or a certain alignment condition** where the results are more effective.
- About that range, there **seems to be a correspondence** throughout all the models, which translates into a certain area of occupancy of a face in the image.
    - **S-MAD binary classification approach** varies between about 50% and 60%
    - **S-MAD fused classification approach** varies between about 42% and 77%
    - **D-MAD fused classification approach** varies between about 56% to 70%

- The **face region** seems to be the dominant activated across all the alignment settings

Scenario	Approach	Alignments	a	b	c	d	e	f	g	h	i	j	k
S-MAD	Binary Classification	AGIR value for bona	4.48	3.35	2.02	1.81	1.59	1.40	1.77	1.53	1.22	1.24	0.95
		Average AGIR value for morphs	1.69	2.02	2.25	2.39	3.32	2.92	2.41	2.52	2.69	2.56	1.79
	Fused Classification	AGIR value for bona	2.35	1.93	2.67	1.39	1.06	1.84	1.09	1.51	1.33	1.43	0.82
		Average AGIR value for morphs	1.44	1.37	1.88	1.04	1.11	1.46	0.78	1.37	0.78	1.01	0.62
D-MAD	Fused Classification	AGIR value for bona	1.87	2.31	1.13	1.81	2.23	1.08	1.97	0.87	1.43	1.20	0.69
		Average AGIR value for morphs	1.83	1.51	1.04	1.42	1.54	1.16	1.48	0.84	1.26	0.91	0.81

- **However**, for both **fused classification** cases (S-MAD and D-MAD), the background seems to have more influence on detection when compared to the S-MAD binary classification approach(lower AGIR values)

- These two fused classification approaches achieve the best performances, which may indicate that **the background of the image does influence the results to some extent.**

FRVT NIST MORPH Benchmark Resuts

Top-performing models
for *fused classification*
approaches
(S-MAD and D-MAD)

→ **Submitted**
(Report 20 June 2023)

Performances evaluated by
comparison
with SOTA MAD approaches.

Test can be divided into
single-image and **differential**
cases

Multiple datasets created
using a diversity of
methodologies.
Tier 1, 2 and 3

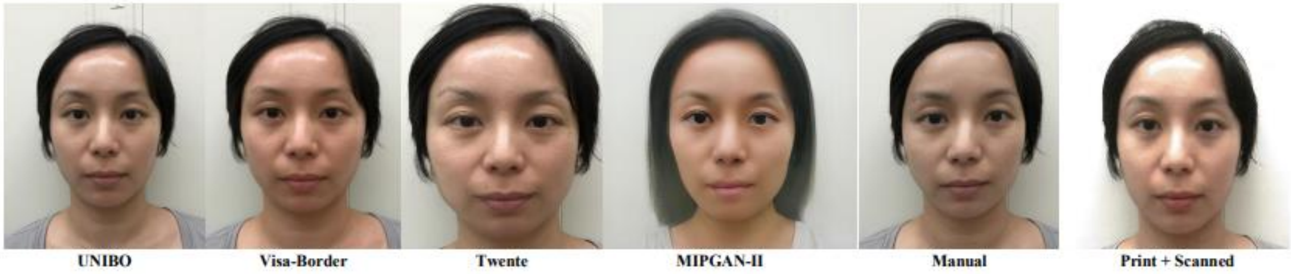

Tier 2- Automated Morphs Analysis

[UNIBO Automatic Morphed Face Generation Tool v1.0 Dataset]

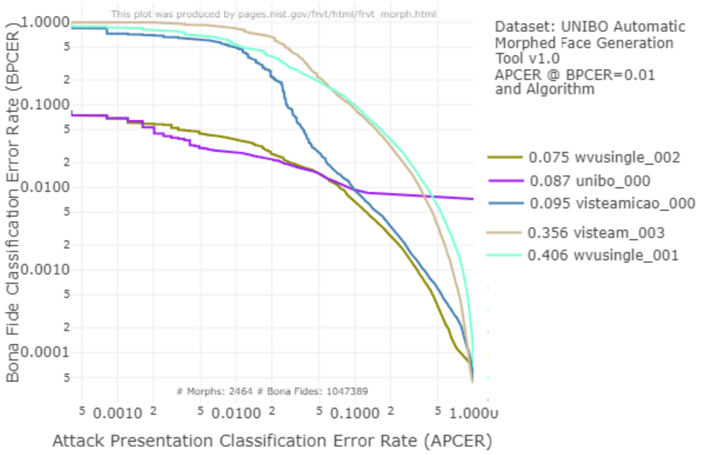

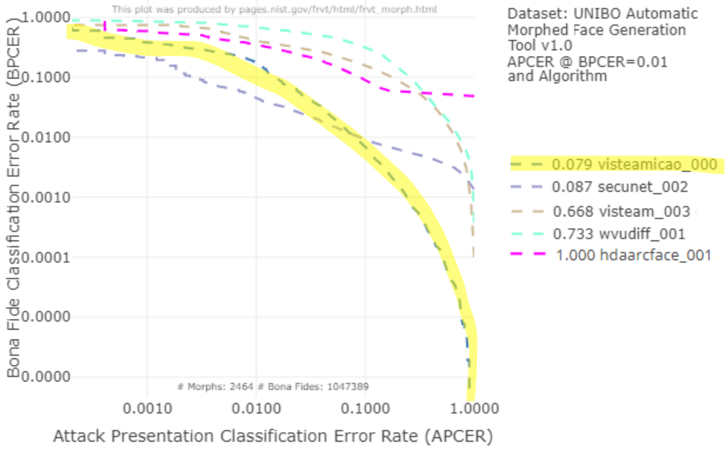

APCER@BPCER=δ							
	Algorithm	δ=0.1	δ=0.01		**Algorithm**	δ=0.1	δ=0.01
Single-image	**wvusingle-002**	0.000	0.075	**Differential**	**visteamicao-000**	0.014	0.079
	unibo-000	0.000	0.087		secunet-002	0.003	0.087
	visteamicao-000	0.027	0.095		visteam-003	0.171	0.668
	visteam-003	0.091	0.356		wvudiff-001	0.257	0.733
	wvusingle-001	0.101	0.406		hdaarcface-001	0.089	1.000

Our model outperforms all others in **the differential case**

Visa-Border Dataset

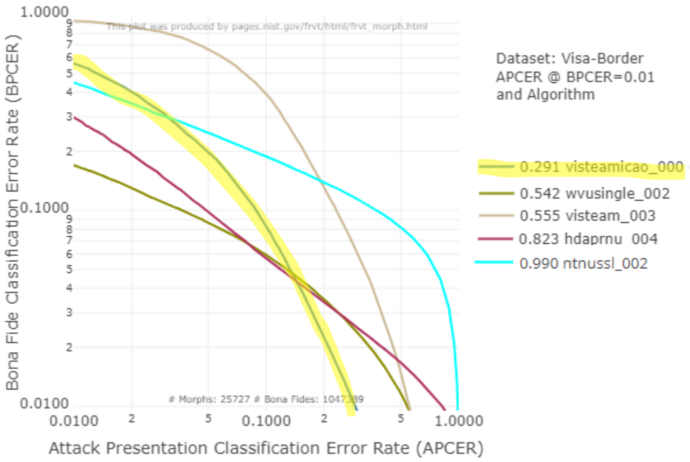

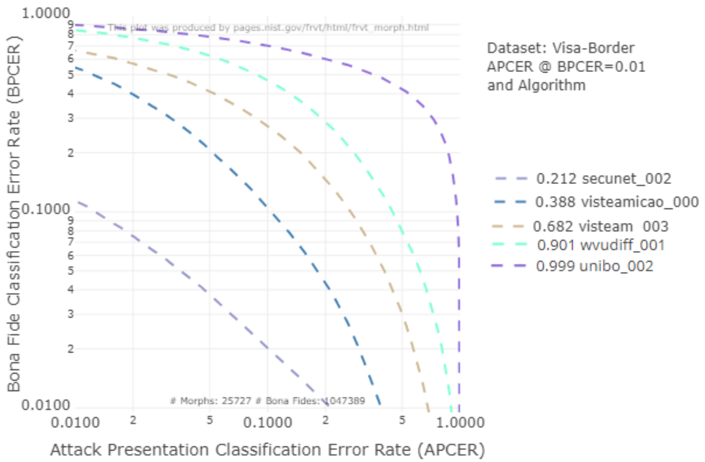

APCER@BPCER=δ							
	Algorithm	δ=0.1	δ=0.01		**Algorithm**	δ =0.1	δ =0.01
Single-image	**visteamicao-000**	0.089	0.291	**Differential**	**secunet-002**	0.013	0.212
	wvusingle-002	0.037	0.542		visteamicao-000	0.105	0.388
	visteam-003	0.232	0.555		visteam-003	0.271	0.682
	hdaprnu-004	0.049	0.823		wvudiff-001	0.447	0.901
	ntnussl-002	0.375	0.990		unibo-002	0.966	0.999

Our model outperforms all others in **the Single case**
(morph miss rate of 0.291 at a false detection rate of 0.01

Tier 3- High Quality Morphs Analysis

Great significance as these datasets closely resemble real-life situations

Manual Dataset – More realistic for **Single Image** case

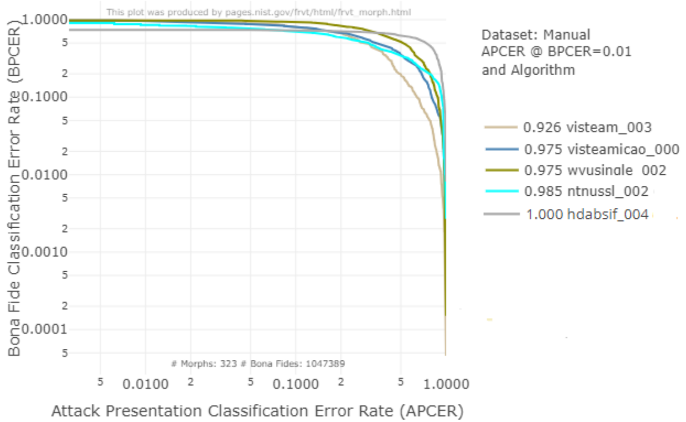

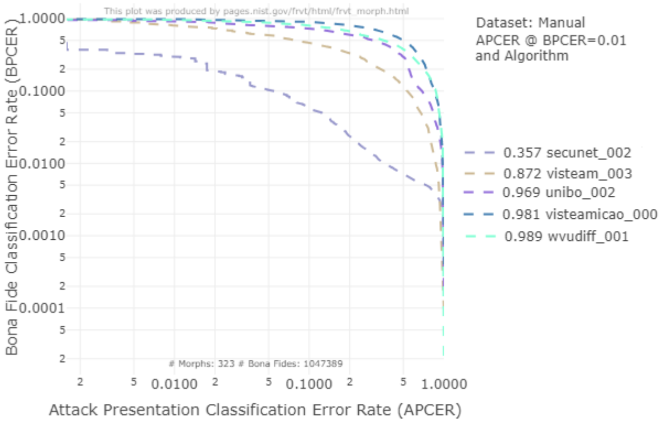

APCER@BPCER=δ							
	Algorithm	δ=0.1	δ=0.01		**Algorithm**	δ=0.1	δ=0.01
Single-image	**visteam-003**	0.641	0.926	**Differential**	**secunet-002**	0.055	0.357
	visteamicao-000	0.802	0.975		visteam-003	0.531	0.872
	wvusingle-002	0.879	0.975		unibo-002	0.689	0.969
	ntnussl-002	0.938	0.985		visteamicao-000	0.853	0.981
	hdabsif-004	0.969	1.000		wvudiff-001	0.873	0.989

Algorithms do not exhibit robust generalization across various unseen morphing techniques.

Print + Scanned Dataset - More realistic for **Differential** case

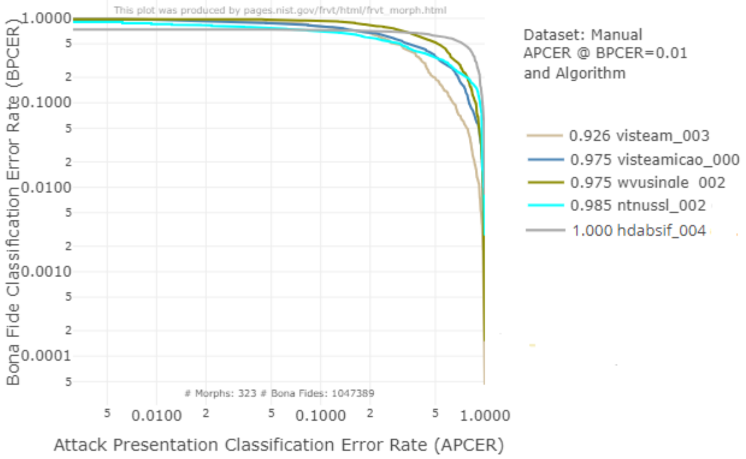

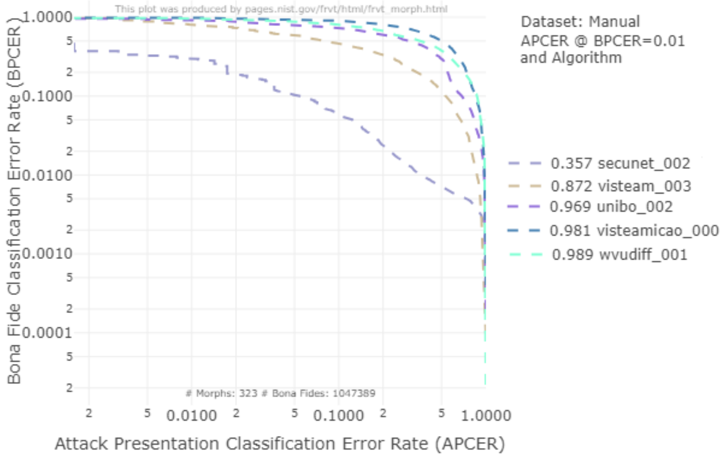

APCER@BPCER=δ							
	Algorithm	δ=0.1	δ=0.01		**Algorithm**	δ=0.1	δ=0.01
Single-image	**wvusingle-001**	0.271	0.721	**Differential**	**secunet-002**	0.012	0.176
	unibo-000	0.420	0.777		unibo-002	0.070	0.280
	visteam-003	0.424	0.788		visteamicao-000	0.426	0.751
	visteamicao-000	0.453	0.819		visteam-003	0.680	0.926
	hdafvdet-001	0.879	0.992		wvudiff-001	0.756	0.953

when compared to other SOTA approaches, our model achieved a **competitive position**.

Conclusions

- In this dissertation, the main goal was to **evaluate the influence of the context of an image in the detection of face morphing attacks.**

- The initial step involved **creating an ICAO-compliant dataset** by combining and pre-processing several datasets.

- Throughout the **different alignment conditions**, the face's occupancy area in the image varies, and consequently, so does the context information.

Conclusions

- Through extensive experiments, **a possible alignment range has been determined** at which Morphing Attack Detection (MAD) is most effective.

- However, the **overall impact** of image context for the face morphing detection **appears to be limited.**

- On **NIST MORPH benchmark**, the results of the presented models demonstrated good performances in several benchmarks. Reaching the state-of-the-art SOTA level in some of them.

Future Work

- **Removing** the background from the image.

- **D-MAD *binary classification*** approach formulation.

- Explore other **explainability tools** or **attention mechanisms** in order to provide more accurate and realistic insights about the decision-making process.

- Generate a dataset that includes **print morphs** in order to train and evaluate the models more realistically, potentially leading to other outcomes.

Thank you!

QUESTIONS ?