Morphed Image Detection using Local Spectrum Analysis

Eitel Yvan Ndeh de Mbah

Abstract

According to research publications, face recognition systems are prone to morphed images attack. To address this issue, many Morphing Detection Attack techniques have been developed. These techniques are classified in three different categories based on the different images features exploited in the technique. One of the biggest challenges encountered when evaluating the performance of a developed technique is the fact that the technique is dependent on the morphed images dataset that was used in developing the technique. This morphing dataset is based on the method used to create the morphed images. This implies the developed technique may perform well on the database that was used to develop the technique, but once a different database is used, the developed method may no longer be able to detect morphed images accurately. In this work, a morphing attack detection (MAD) technique based on local spectrum analysis is presented. In this technique the absolute total power content over the spectrum of the images is used to differentiate between morphs and genuine images (mostly referred to as bona fides in biometrics). The analysis is made locally by splitting each facial images into six facial features and the classification based on each feature is attributed an individual binary score. This score refers to a right or wrong classification. To detect if a given image is a morphed image, these scores are then summed to obtain a total score and based on this score the proposed algorithm can determine if the image is morphed or not. FRGC database is used to train, validate and test the performance of the algorithm. A second database, an AMSL database is used to investigate the robustness of the proposed algorithm. Experimental analysis showed an ACER of 1.59% when the technique was tested on the FRGC database and an ACER of 38.24% on the AMSL database. ACER represents an error rate that was used to evaluate the performance of the technique on moprhed and genuine images.

Keywords—*Bona fide, Morph, Score*

1 Introduction

Face images can be used in the production of identification documents and in digital security systems. In some countries, for issuing identification documents like passports and ID cards, the applicants are usually required to provide a picture of their face in the required format [1]. The applicant can apply some changes and modifications to this picture before submitting it to the authority in charge of producing the document. One of the changes that can be applied to such a picture is called morphing and the result of such changes is called a morphed image.

A morphed image can be defined as the result of a morphing process in which two or more images are combined to form a new image. Depending on the morphing procedure used, this new picture can easily be used by any of the people whose image contributed to the morph for identity verification purpose. This is because the new picture will carry the facial traits of both people and will make it even difficult for the human eye to notice that the image has been digitally manipulated. If such a picture is used in a passport application process, the passport that will be produced may deceive the ABC gates at airports or even the immigration officer at the immigration control during matching. Matching in the context of face images and face recognition systems represents a process whereby a live captured face image of a person or any other image of that person is compared to a stored image of the person in a database for identity verification purpose. Currently there exists some free and payed image processing and modification software available to the public which makes it easier for people to modify their pictures and images. Some of these software include OpenCV/dlib, FaceMorpher, FaceFusion, GNU Image Manipulation Program v2.8 (GIMP), etc...[1][2][3]. These applications, all give the possibility to morph two or more facial images together to form one a new facial image called a morphed image. There are mainly two types of morphed images known as "Complete" morphs and "Combined/Splicing" morphs. To obtain a complete morph, the two genuine images involved are each partitioned into a triangular mesh with nodes at certain spots, and the triangles obtained in both images are warped and blended together to fit an average geometry which results in the morphed image [4]. Complete morphs usually suffer from ghosting artifacts in the hair region which makes it apparent even to the naked eye that the image is morphed as shown in Fig. 1c. To compensate for this limitation in complete morphs, a different type of morph can be obtained which has no ghosting artifacts but can only match to one of the person involved in the morphing process perfectly. This is usually referred to as a spliced morph, since the face geometry of the morphed image will be inherited just from one person's image, hence matching only that person. Fig. 2 and 3 show two examples of spliced morphs.

(a) Person A (b) Person B (c) Complete morph

Figure 1: Original images and resulting complete morph

(a) Person A (b) Person B (c) Spliced morph

Figure 2: Original images and resulting Spliced morph based on person A's image

(a) Person A (b) Person B (c) Spliced morph

Figure 3: Original images and resulting Spliced morph based on person B's image

Face recognition systems may therefore be vulnerable to morphing attacks and can pose a security risk. This paper interest is therefore on Morphing Attack Detection (MAD) systems. Some MAD approaches have been published, see section 2. These proposed approaches can be categorized under two scenarios;

1. *No-reference MAD:* With this approach, the detector uses a single image to do the analysis and determine whether the image under check is a morph or not.
2. *Differential MAD:* In this second scenario, the detector still uses the single image plus an additional trusted live captured image of the owner of the single picture to do the morphing check.

In this study, the approach that was followed was a No-reference MAD. Under this scenario, a morphing algorithm based on spectral analysis was developed and evaluated on two types of image databases. The facial images were analysed locally by using six chosen facial key features to differentiate morphed images from bona fides.
From the related study, it is observed that one of the main challenge in this research field is the data set used in the methodology. This is because the MAD technique developed is dependent on the data set and as such it has not yet been made possible to generalise any of the techniques for different data sets. Hence we would like to investigate the influence of the data set on our developed MAD technique. By so doing, I would like to answer the research question "To what extent can we use local spectrum analysis for Non-reference morphing attack detection?".

The remainder of this paper is organised as follows: Section 2 will give an overview of some developed MAD techniques. In Section 3, the proposed morph detection method and its implementation is described. Section 4 shows the experimental results obtained and gives an answer to the research question. Finally a conclusion is presented in section 5.

2 Related Work

The literature overview focuses more on the No-reference (blind) MADs. There are three categories under which the existing face morphing detection methods can be classified. This includes "micro-textures feature based methods", "JPEG compression feature based methods" and "deep learning feature based methods" [2].
Micro-features based methods: uses micro traces in the image that usually arise from the computer operations and transformations that were made on the original image to obtain the morphed image. [5] gives an example of a technique that was developed based on this method. The micro trace that was exploited here was the sensor pattern noise (SPN). This is a noise pattern that is generated and added to any image taken by a camera and it is unique to every camera. The SPN of any image directly obtained from a camera often has a high frequency energy content meanwhile that of a morphed image does not. This is because during the morphing process this energy is evenly averaged by the different morphing operations. The resulting SPN of a morphed image is therefore seriously affected and this difference can be used to distinguish between genuine and morphed images. A facial statistical quantization feature is obtained from the SPN of the image under evaluation and is fed to a linear SVM classifier to determine if the image is morphed or not. This method saw a higher performance in complete morphs detection than splicing morphs but yet still had a better performance for splicing morphs compared to the mentioned techniques in the literature [5]. A further experiment was carried out on JPEG compressed images and the method showed a performance that was proportional to the quality factor (QF) of the JPEG compressed image with the best performance achieved with QF greater than 90.
JPEG compression feature based methods: As the name implies it is a method that uses the equal compression properties of the JPEG format of an image to differentiate between morphs and bona fides. [4] is a good example of a technique that was developed based on this method. The Benford features of images were plotted, and it was observed that images that were compressed only once followed a logarithmic distribution. Since the morphing process takes as input an image (in this case a JPEG compressed image) and gives as output a morphed image that is again saved in JPEG format. This implies the output image will now contain a double JPEG compression. If the Benford features of such images is plotted, a deviation from a logarithmic distribution is observed because of the double compression. This deviation was therefore extracted and used by the proposed method to detect morphed images. The fitting error of the original image to the logarithmic distribution according to [4] is expected to be higher than that of a morphed image. The fitting error (MSE) and two other parameters were used to classify original and morphed images again with the hep of a linear SVM. One remark in this study is that the dataset of morphed images used in the experiments had the constraint that the two persons involved

in the morphing process should have the same sex and ethnicity.
The last technique presented in the available literature can be considered as micro texture feature based technique. This method is based on PRNU analysis [2] . PRNU is based on sensor noise. It is similar to SPN method but instead of extracting the sensor pattern noise, PRNU values are extracted in spectral and spatial domain and are used for classification. Since this is still an inherited property of an original image, the morphing process will alter the PRNU values due to the warping and averaging operations involved. Compared to the SPN method which only did it analysis in the spectral domain, this method makes use of both the spectral and spatial domain by extracting PRNU values of the images under analysis in both domains. From each of these domains, two features are retained (which simplifies the method) and are fed to a linear SVM for morph detection. It should also be noted that this method uses a local spectrum analysis since it divides the image into cells and each individual cell contributes a feature.
One major observation on each of these techniques is that their performance is mainly dependent on the dataset used for training the SVM i.e, how both the original and morphed images were obtained. If a different dataset (images with different characteristics) is used to test the detector some discrepancies can be observed. Cross database evaluation was only observed in the proposed PRNU method and the best performance was achieved with an average equal error rate of 11.2%. It was proposed that a sophisticated approach based on machine learning techniques could therefore be used with multiple PRNU features to improve the performance of the method on different databases.

3 Methodology

The developed algorithm is based on spectral differences resulting from the warping procedure during the morphing process. Additionally, the algorithm attempts a local analysis of the images by segmenting the image into six facial features and at the end combines the six different facial features analysis to generate a final decision on the image class (morph or bona fide). The facial features extraction was motivated by the guided face predictor [6] which could be used to extract specific features of the face and determine how important is each facial feature in the morphing detection. An arbitrary facial extraction could also be considered for a later analysis. The developed technique consists of five main steps which are illustrated in Fig. 4. These various steps and consequently the experiments that were conducted following these steps are discussed in the remainder of this section in more detail.

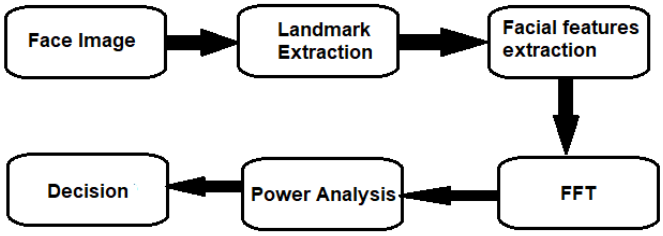

Figure 4: Processing steps of the proposed local spectrum based morphing detection technique

3.1 Databases used

The first step was to choose the databases on which the method would be trained and validated. Two different datasets were used to develop and validate the algorithm both from an FRGC database. The first dataset called the training set was constructed and only used to obtain an estimate of the thresholds range that could be used to correctly classify images based on the extracted absolute total power content of the images. The second dataset labelled as the validation dataset was then used to select the exact parameters to use for each feature classifier as will be described in the later section. Then another set from an FRGC database was used as a test dataset to evaluate the performance of the technique on images from the same database (FRGC) .
Lastly a test dataset obtained from an AMSL database was used to test the algorithm and to investigate its robustness on a different database. The FRGC database was extracted from the Face Recognition Grand Challenge dataset [7] and the AMSL database from [8],[9] and [10].

3.2 Facial Features Extraction

Since the method aims to apply local spectrum analysis on the facial images, an adopted scheme to perform these local analysis is to split the face image into six facial features as follows;

1. The eyebrows
2. The left eye
3. The right eye
4. The nose
5. The mouth
6. The chin with some skin texture

To be able to do this, a shape predictor was used to determine 68 points on the face called facial landmarks [6]. Fig. 5a shows a face image and the corresponding 68 points obtained are displayed as red dots on the face image as shown in Fig. 5b. Then using these points it was possible to extract each facial feature as an approximated area covered by chosen reference points. This particular splitting of features was chosen to also investigate the influence of each facial key features on the MAD detection technique. Another implementation of the local analysis could be to make a random splitting of the facial image into three rectangular boxes or four quadrants and analyse each quadrant separately.

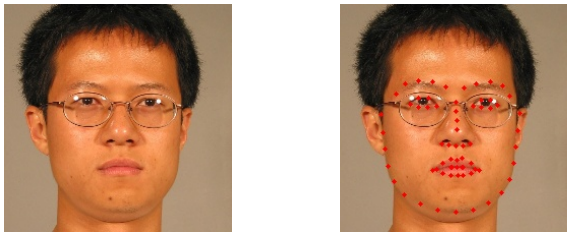

(a) Original (b) landmark coordinates

Figure 5: Image and corresponding 68 points coordinates facial landmarks

The features extraction was implemented using a calculated distance between appropriate coordinates for each feature. This resulted into facial features with different dimensions for each image. For uniformity in the dimensions of the extracted feature, after extraction, they were resized. The resized dimensions were fully based on observation after running a couple of tests on different images motivated by the intention to keep just the feature under consideration in the resulting image. The adopted dimension for each feature is shown in table 1 and the extracted facial features for the facial image in Fig 5 is shown in Fig. 6.

Table 1: Sizes of the different extracted facial features

Facial features	**Sizes**
Eyebrows	25×120
Left eye	25×50
Right eye	25×50
Mouth	40×80
Nose	60×45
Chin	25×80

.

Figure 6: Facial features extracted; eyebrows, left eye, right eye, mouth, nose, chin

3.3 Fast Fourier Transform (FFT)

An image is represented by an array of pixels in spatial domain. For spectral analysis, the image is converted to frequency domain by means of the Discrete Fourier Transform (DFT) since the images are not infinite in spatial domain. The DFT represents the Fourier Transform of a finite sequence (which in this case will be our image). To efficiently compute the DFT of an image as described in [11], Fast Fourier Transform is used and the result is a magnitude spectrum. Magnitude spectrum is a representation of the frequency components of the image in spectral domain. Before applying the FFT operation, the image under analysis is first converted to grey-scale to have a 2D array. This is because the algorithm is developed in Matlab and the FFT operation is limited to 2D arrays. Another possibility would be to apply the FFT on each of the RGB channels of the image under analysis. This would have probably been more effective since all the frequency information content of the image is preserved but for simplicity the the grey-scale conversion was used. An example of a facial feature image (eyebrows) and the resulting magnitude spectrum is shown in Fig. 7.

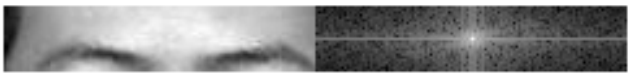

Figure 7: Image of eyebrows in grey-scale and corresponding magnitude spectrum

3.4 Power Analysis

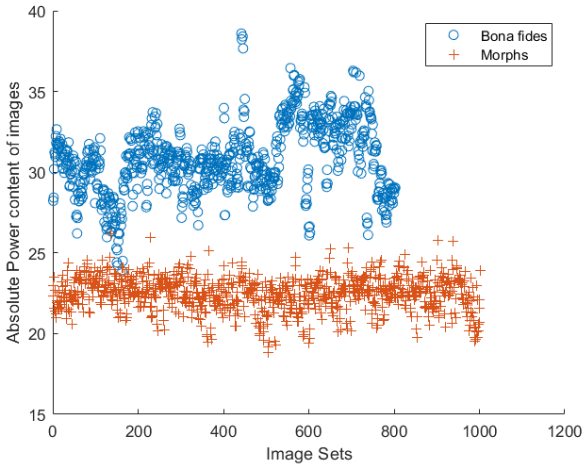

Figure 8: Scatter plot of power content for the eyebrows of the different images in the training set

It is important to note that the FFT result is in the complex domain which originates from the definition of Fourier Transforms. After obtaining the frequency content of the image, the spectral features that could be used for analysing the nature

of the image (morph or non-morph) had to be determined. Using the power spectral density estimation, it's possible to estimate the power contribution over the selected frequencies. Since the dimensions of my features under analysis are not large, I simplified the analysis to the whole image of the feature thereby computing the absolute total power content of the facial feature over all the frequencies. Moreover, since the image is finite, the absolute total power can be obtained using equation 1. This equation was motivated and derived from [11].

$$P_{img} = \frac{1}{MN} \sum_{i=0}^{N-1} \sum_{i=0}^{M-1} (|X[k] \times \bar{X}[k]|)(|Y[k] \times \bar{Y}[k]|) \quad (1)$$

Where $X[k]$ and $Y[k]$ represent the Fourier transform of the image in the x and y-coordinate respectively. $\bar{X}[k]$ and $\bar{Y}[k]$ represent the complex conjugate of $X[k]$ and $Y[k]$ respectively. Lastly, the product MN is the total number of pixels in the image.

For our analysis the sum of the of the absolute power content over all frequencies of the facial feature is computed which is essentially a finite impulse response filter. It is expected that for a Bona fide, this absolute total power should be larger than that of a morph. During the morphing process, the two images are usually added together and averaged to obtain the morph. This averaging step when translated in the frequency domain will reduce the DFT coefficients and hence the power content also expected to be lowered. Moreover the high frequencies that could be observed in the bona fides might be lowered in the morphs since this high transitions would have also been averaged and smoothened during the morphing process. Therefore the information content at higher frequencies will be higher for bona fides than for morphs, so computing the absolute power helped to quantify this information content. With this in mind, using the training dataset, the absolute total power of each facial feature was computed for each image and the result plotted to observe if there will be an exploitable difference in the power content of genuine and morphed images.

In this analysis, the absolute total power was chosen to be the spectral feature to observe but this was not the only feature that could be exploited. The sum of the sum of the magnitudes of each frequency could also be extracted and compared for both morphs and genuine images. But for 2D signal processing, the power spectral density gives a better estimation of the frequency content since the multiplication step takes care of the complex domain of the Fourier transform without adding or reducing any information.

Figure 8 shows the computed absolute total power of eyebrows for the images in the training set. From this plot, we could observe that there is indeed a difference in the power content of bona fides compared to that of morphed images. The magnitude of this calculated power wasn't of interest but instead the magnitude of the minimum power content of bona fides and the magnitude of the maximum power content of morphed images were the observed quantities. With these two values, a threshold range that could be used to classify bona fides from morphed images using a specific facial feature was determined. Using the validation dataset and the extrapolated threshold range for each feature, a ROC curve was plotted to observe the performance of the classifier with different thresholds as shown in Fig. 10. From the ROC curve, the threshold at which a False Acceptance Rate of 0.1 % is achieved was extrapolated. An FAR of 0.1% used as a criterion to adopt a threshold was motivated by the work in [12]. This rate was used to optimise three-dimensional face recognition and gave a good result. This threshold was then adopted as the final threshold with which a particular facial feature of an image will be used to classify it as a bona fide or a morph based just on that feature.

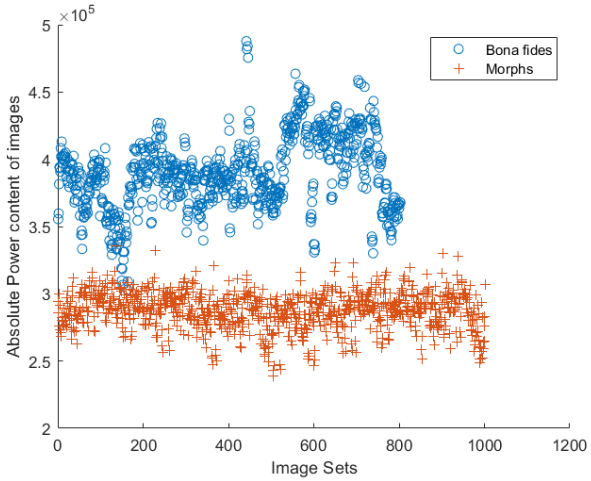

Figure 9: Scatter plot of power content for the eyebrows of the different images in the training set

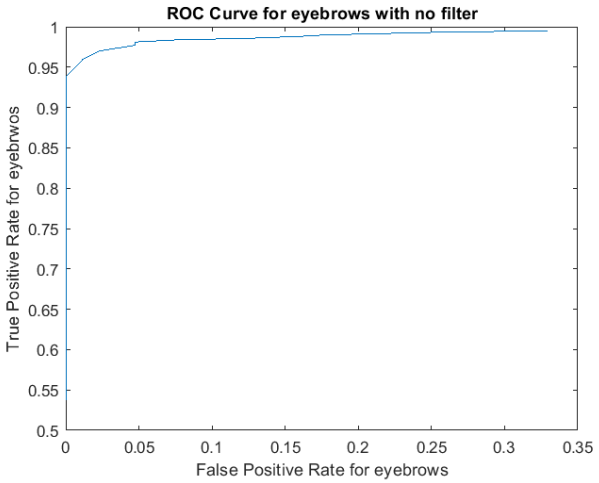

Figure 10: ROC for the eyebrows using the validation dataset

An intermediate step between the FFT and the Power analysis was later introduced to observe the absolute total power contributed by higher frequencies. The computed DFT was passed through a high pass filter and the total absolute

power of the filtered DFT determined. This was to see if focusing just at the power content of higher frequencies of an image (facial feature) could give a better detection performance. The high pass filter was implemented by generating a Matlab circular filter that will be able to block the undesired low frequencies power content. Depending on the facial feature dimension under consideration, different radii of this filter was used. Fig. 9 shows a similar plot as that in Fig. 8 but now with the DFT filtered. It was observed for this plot that the magnitude of the power highly increased which was as a result of the filtering process but the absolute difference between the bona fides power content and morphs power content was maintained.

3.5 Decision

Our proposed morph detection technique is made up of six classifiers from the six different facial features extracted. The six decisions then needed to be combined to one single decision to classify the entire facial image. To do this, a scoring system is used on the validation dataset whereby a score of "1" is attributed to every positive decision and a "0" to every negative decision. First, all the morphs were classified for each facial feature using the corresponding adopted absolute power threshold. Then for each image that was rightly classified as a morph (True Positive), a score of 1 was attributed to this image and for the images that were wrongly classified as bona fides (False Negative), a score of 0 was given. This was repeated on the set of bona fides and for every rightly classified bona fide, a score of 0 was given (True Negative) and for every wrongly classified bona fide a score of "1" was recorded.
After obtaining the six different scores with respect to the six different facial features, the scores were added for each image giving a maximum score of "6" and a minimum score of "0" for each image. These summed scores were plotted as two normalised histograms as shown in Fig. 11 for all the images in the validation set. One histogram representing the distribution of scores for morphs and the other displaying the distribution of scores for bona fides. The two charts overlap between scores 3 and 5.

1. With a score of "3", there is roughly a probability of 0.05 of wrongly classifying a bona fide as morph and a probability of less than 0.01 to wrongly classify a morph as a bona fide.
2. With a score of "4", there is probability of 0.04 of wrongly classifying a morph and a 1% probability of wrongly classifying a bona fide
3. Lastly, a score of 5 will give a 22% probability of wrongly classifying a morph and a probability of less than 0.01 of wrongly classifying a bona fide.

Therefore score 4, gives a more accurate classification with a lower error rate for both morph detection and bona fides identification. This was then accepted as the final score to classify a given facial image as either a morph or bona fide for the developed MAD technique.

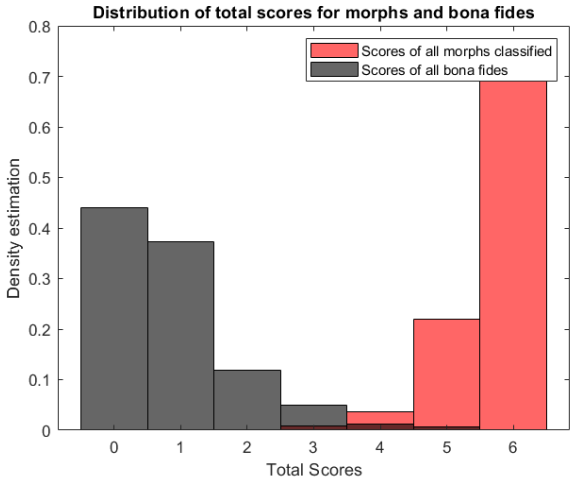

Figure 11: Normalised Distribution of total scores for morphs and bona fides

Table 2 gives the performance of the algorithm on the validation dataset using the decision score of 4.

Table 2: Performance of algorithm on validation dataset

Subjects	**Total**	**Rightly Classified**	**Wrongly Classified**	**Classification Error**
Bona fide	1090	1068	22	2.02%
Morphs	1477	1464	13	0.88%

The same analysis was repeated on the filtered DFT's and the results uploaded in the appendix. At the end of the validation process, the radius and associated threshold for each facial feature was selected based on the combination which gave the highest classification performance with the validation dataset. Table **??** gives a summary of all the parameters that were then adopted for the algorithm and that will be used in the next section on the test dataset for performance evaluation of the technique.

Table 3: Performance of algorithm on validation dataset

Feature	**Filter Radius**	**Power threshold**
Eyebrows	No filter	30
Left Eye	No filter	27
Right Eye	No filter	28.75
Nose	No filter	25.3
Mouth	No filter	27.5
Chin	No filter	19.05

4 Results and Discussion

As previously mentioned in the methodology, the developed MAD technique was tested on two images datasets to investigate its robustness on datasets. [5] presents three standadized

ISO metrics used to evaluate the overall detection performance of an MAD technique. These metrics include: *Attack Presentation Classification Error Rate (APCER)*, *Bona Fide Presentation Classification Error Rate (BPCER)* and lastly the *Average Classification Error Rate (ACER)*. Equations 2, 3 and 4 gives a mathematical definition of these metrics.

$$APCER = \frac{\#\text{ of morphs wrongly classified}}{\text{Total number of morphs in the dataset}} \quad (2)$$

$$BPCER = \frac{\#\text{ of bona fides wrongly classified}}{\text{Total number of bona fides}} \quad (3)$$

$$ACER = \frac{\text{APCER} + \text{BPCER}}{2} \quad (4)$$

4.1 Testing developed MAD technique on FRGC database

A total of 2648 images from an FRGC database were used to test the algorithm. Table 4 below gives a summary of the results obtained.

Table 4: Results obtained with FRGC test dataset

Subjects	Total	Rightly Classified	Wrongly Classified	APCER	BPCER	ACER
Bona fide	1050	1018	32	0.125%	3.05%	1.59%
Morphs	1600	1598	2			

4.2 Testing developed MAD technique on AMSL database

Unfortunately the available database had just 204 bona fides. So there wasn't an extensive test of the algorithm on bona fides as was done on the FRGC database. 2175 morphed images from the AMSL database were used to test the developed MAD technique. Table 5 shows a summary of the results obtained.

Table 5: Results obtained with AMSL test dataset

Subjects	Total	Rightly Classified	Wrongly Classified	APCER	BPCER	ACER
Bona fide	204	49	155	0.5%	75.98%	38.24%
Morphs	2175	2164	11			

From the experimental results and analysis, the proposed MAD method can achieve good detection performance on an FRGC database.But, on an AMSL databse, the method turns to classify almost 76% of bona fides as morphs leading to a very low over all performance of the method on the AMSL database. This implies the thresholds that were used in the validation section were very good for morph detection of both databases. The false classification rate of moprhs in the AMSL database was even lower than 1% which is really good just for morph detection but not that good for the overall performance of the method. With these results, an answer to the research question could be formulated as follows; Local spectrum analysis can be used for Non-reference morph attack detection and the developed method in this study showed an overall error rate of less than 2% for both morphs and bona fides . But the method is still very database dependent and will need to be trained on the database on which it will be implemented for it to perform as desired. Since only one different database was used to test the method, it can not yet be concluded that the method won't work on other databases. The method performance could not be compared to other published methods because the databases used were different. The intermediate step taht was introduced in which the DFT was filtered din't give better results. This could be because the filtering technique employed (using circles) was not efficient. Another filtering technique that could be implemented to investigate the power content at high frequencies could be using rectangular blocking filters instead of circular filters. Also the sum of the magnitudes of the frequencies was investigated using the training database but there was no observable threshold with which we could be differentiate morphs from bona fides with the same or higher performance as that obtained with the absolute total power analysis.

5 Conclusion

Face recognition systems can be exposed to morphing attacks. In this work, an overview of some of the developed MAD algorithms was looked at and from there a no-reference face morph detection algorithm based on local spectrum analysis was developed. The morphing process manipulates the frequency content of an image and in most cases the resulting frequency content is averaged and lowered which is usually not the case for genuine images. With this it was expected that the absolute total power content of morphs would be lower when compared to that of genuine images. This difference was then used to for morph image detection while carrying the analysis on six different facial key features which resulted in six independent sub-classifiers. Experimental results shows that MAD can be done using local spectrum analysis and an ACER of 1.59% observed when tested on an FRGC database. On a different database, the AMSL database a lower performance was observed with an ACER of 38.24% which showed a limitation of the developed method on databases. In a future work, the method could be tested on more databases to see if there could be more than one database for which such a MAD method could work. Then training and validating the methodology on two to three databases could further be investigated to make the algorithm more robust to new databases. Different techniques of analysing the face locally were also proposed and could be investigated.

6 References

[1] M. Ferrara, A. Franco, and D. Maltoni, "The magic passport,"

[2] U. Scherhag, L. Debiasi, C. Rathgeb, C. Busch, and A. Uhl, "Detection of face morphing attacks based on prnu analysis."

[3] U. Scherhag, C. Rathgeb, J. Merkle, R. Breithaupt, and C. Busch, "Face recognition systems under morphing attacks: A survey," *IEEE Access*, vol. 7, Mar. 2019.

[4] A. Makrushin, C. Kraetzer, T. Neubert, and J. Dittmann, *Generalized benford's law for blind detection of morphed images*, Jun. 2018. DOI: 10.1145/3206004.3206018.

[5] L.-B. Zhang, F. Peng, and M. Long, "Face morphing detection using fourier spectrum of sensor pattern noise,"

[6] A. Rosebrock. (2017). "Facial landmarks with dlib, opencv, and python," [Online]. Available: https://www.pyimagesearch.com/2017/04/03/facial-landmarks-dlib-opencv-python.

[7] (2010). "Face recognition grand challenge (frgc)," [Online]. Available: https://www.nist.gov/programs-projects/face-recognition-grand-challenge-frgc.

[8] F. R. L. L. Set. (), [Online]. Available: https://figshare.com/articles/Face_Research_Lab_London_Set/5047666.

[9] U. E. Dataset. (), [Online]. Available: http://pics.stir.ac.uk/.

[10] (), [Online]. Available: https://omen.cs.uni-magdeburg.de/disclaimer/index.php.

[11] R. Veldhuis. (2019). "Lecture notes discrete-time signal processing," [Online]. Available: https://canvas.utwente.nl/courses/8102/pages/lectures?module_item_id=225713.

[12] L. Spreeuwers, "Breaking the 99% barrier: Optimisation of three-dimensional face recognition," *IET Biometrics*, 2014. DOI: 10.1049/iet-bmt.2014.0017.

A Filtered Scatter plots of training set

A.1 scatter plots for the eyebrows

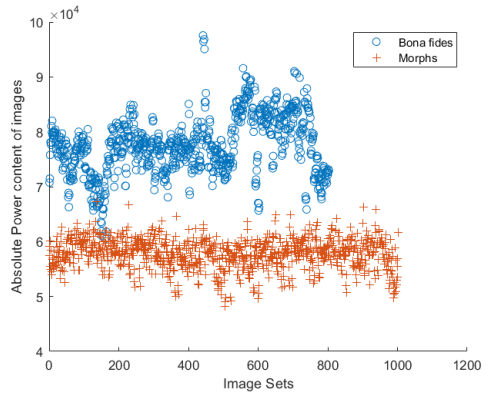

(a) Scatter plot of power content of eyebrows for filter radius 5

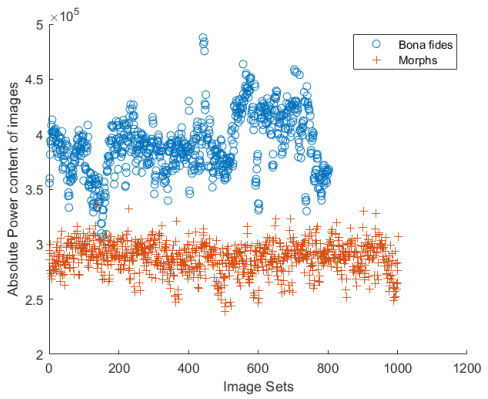

(b) Scatter plot of power content of eyebrows for filter radius 10

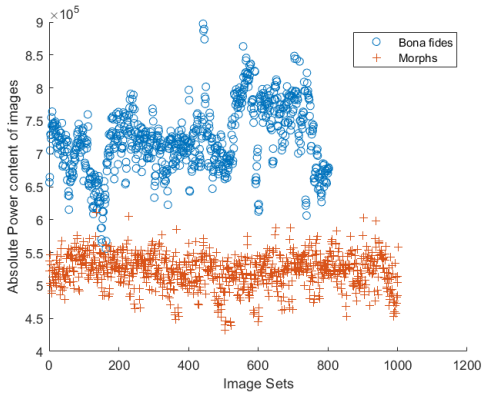

(c) Scatter plot of power content of eyebrows for filter radius 15

Figure 12: scatter plots of the power for the filtered DFT of the eyebrows

A.2 scatter plots for the left eye

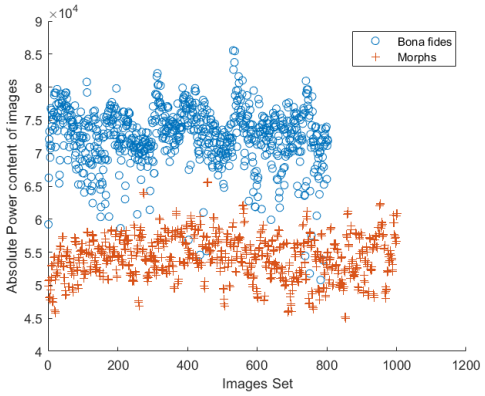

(a) filter radius 5

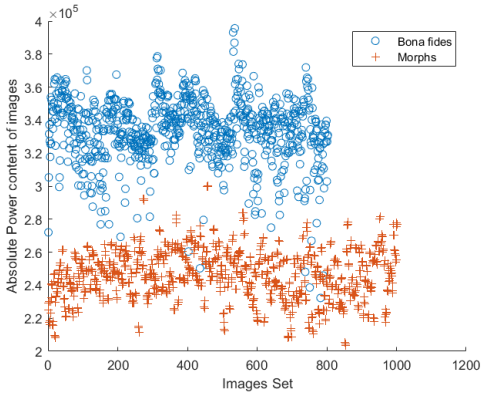

(b) filter radius 10

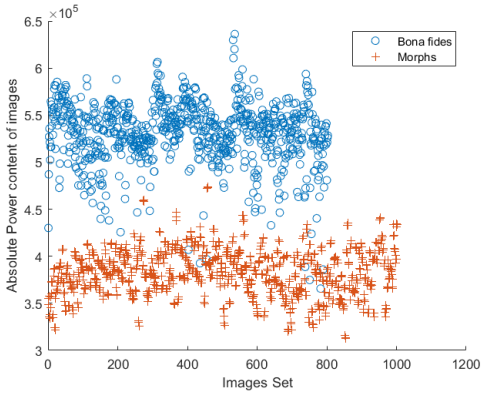

(c) filter radius 15

Figure 13: scatter plots of the power for the filtered DFT of the left eye

A.3 scatter plots for the right eye

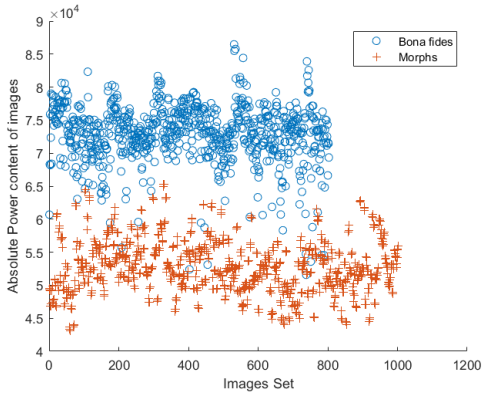

(a) filter radius 5

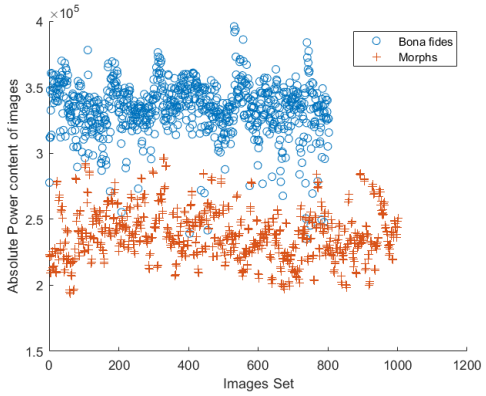

(b) filter radius 10

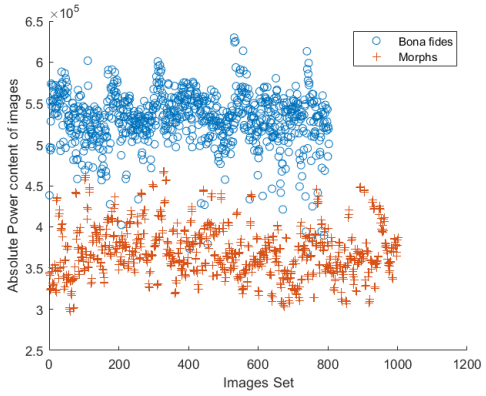

(c) filter radius 15

Figure 14: scatter plots of the power for the filtered DFT of the right eye

A.4 scatter plots for the nose

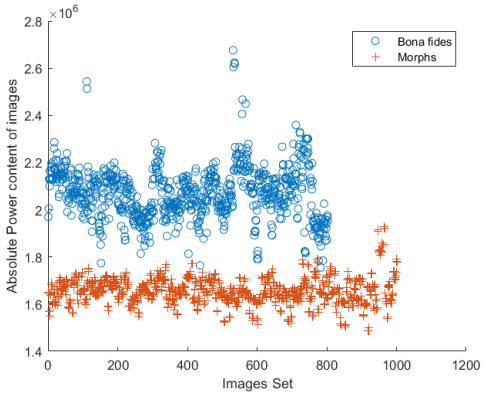

(a) filter radius 20

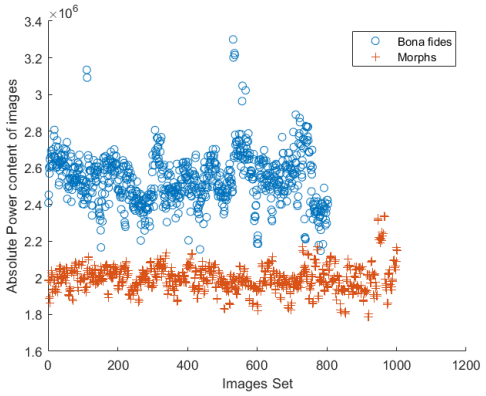

(b) filter radius 25

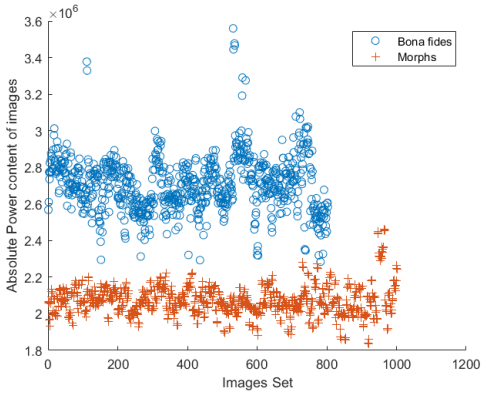

(c) filter radius 30

Figure 15: scatter plots of the power for the filtered DFT of the nose

A.5 scatter plots for the mouth

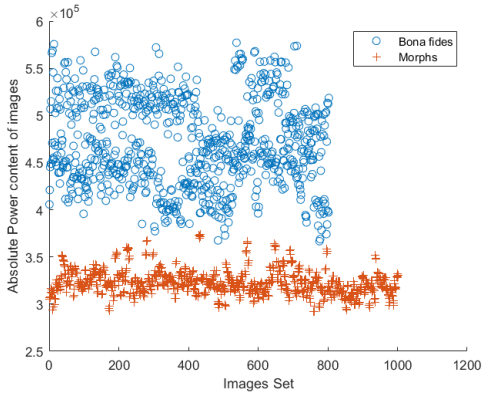

(a) filter radius 10

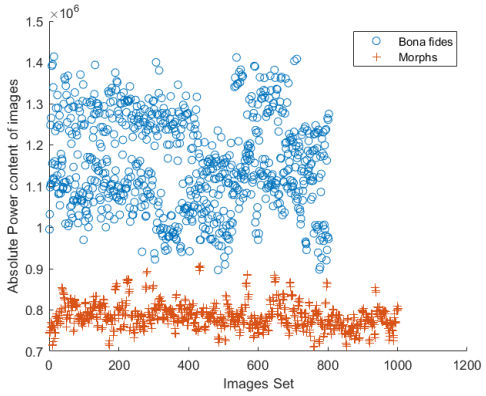

(b) filter radius 15

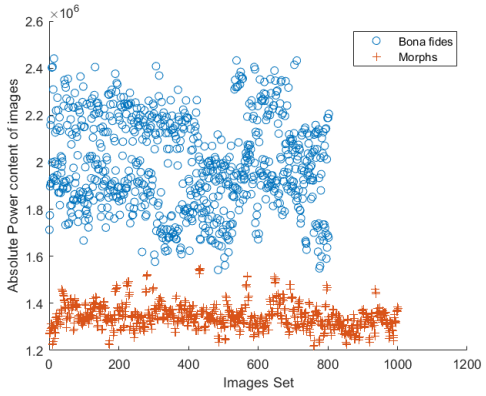

(c) filter radius 20

Figure 16: scatter plots of the power for the filtered DFT of the mouth

A.6 scatter plots for the chin

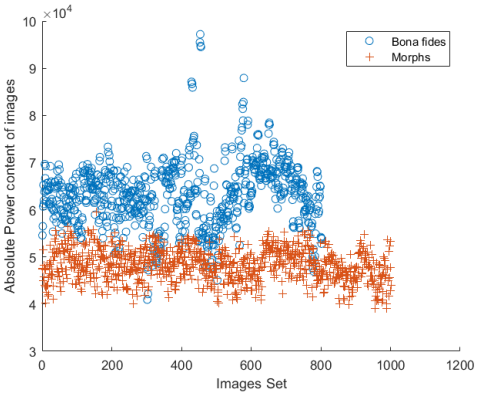

(a) filter radius 5

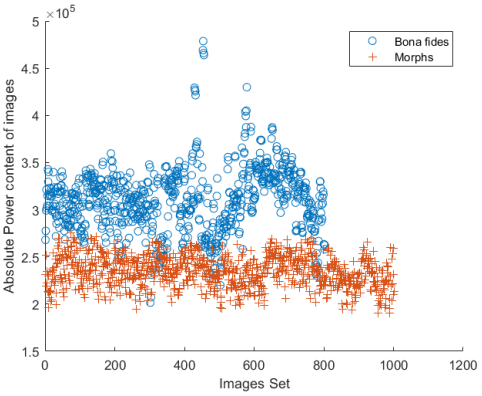

(b) filter radius 10

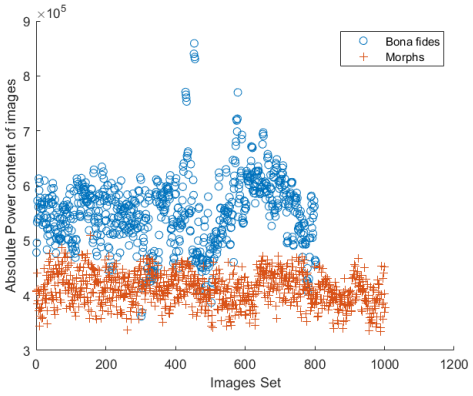

(c) filter radius 15

Figure 17: scatter plots of the power for the filtered DFT of the chin

B ROC curve plots of validation set

B.1 ROC curves for the eyebrows

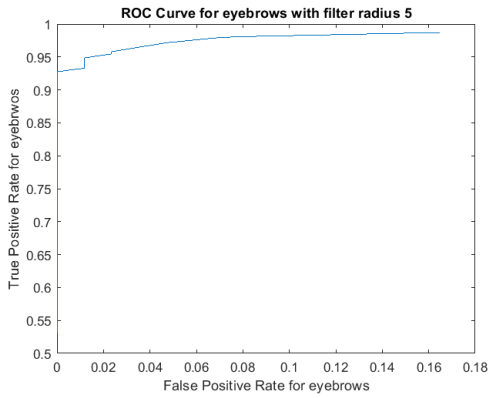

(a) ROC curve for eyebrows for filter radius 5

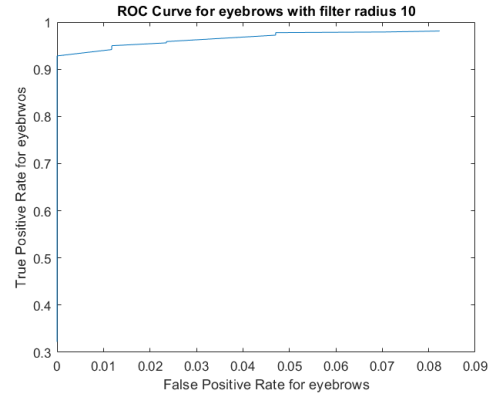

(b) ROC curve for eyebrows for filter radius 10

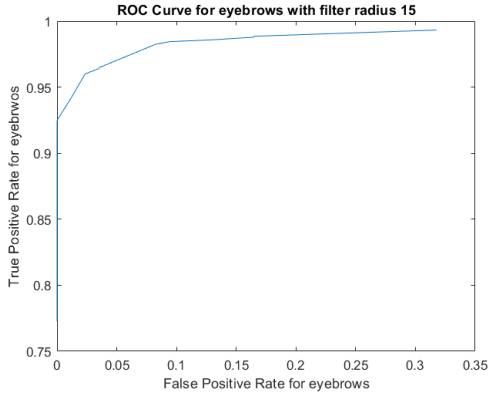

(c) ROC curve for eyebrows for filter radius 15

Figure 18: ROC curves for the filtered DFT of the eyebrows

B.2 ROC curves for the left eyes

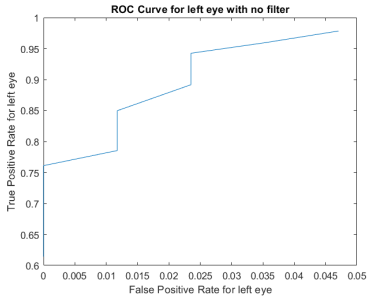

(a) ROC curve for left eye with no filter

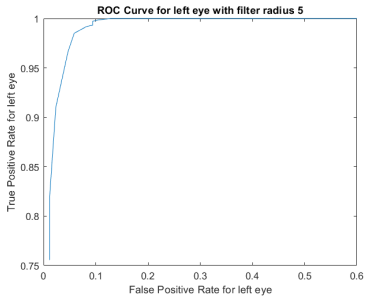

(b) ROC curve for left eye for filter radius 5

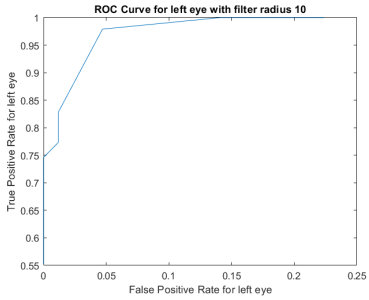

(c) ROC curve for left eye for filter radius 10

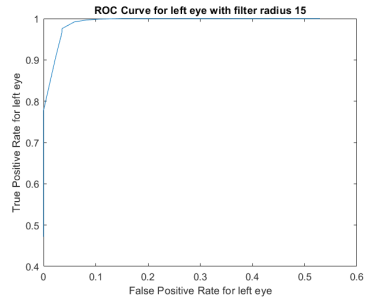

(d) ROC curve for left eye for filter radius 15

Figure 19: ROC curves for the left eye

B.3 ROC curves for the right eyes

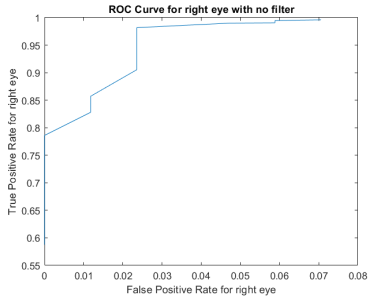

(a) ROC curve for right eye with no filter

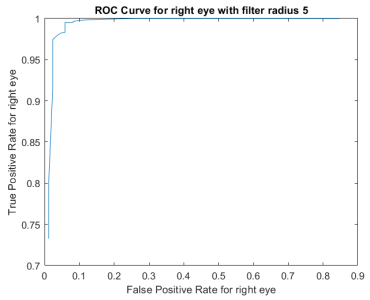

(b) ROC curve for right eye for filter radius 5

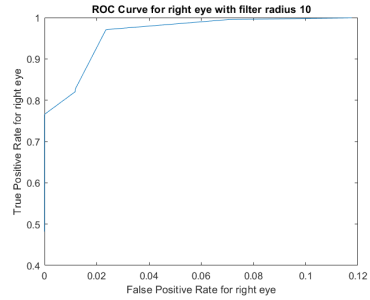

(c) ROC curve for right eye for filter radius 10

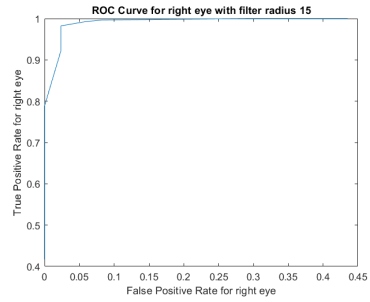

(d) ROC curve for right eye for filter radius 15

Figure 20: ROC curves for the filtered DFT of the right eye

B.4 ROC curves for the nose

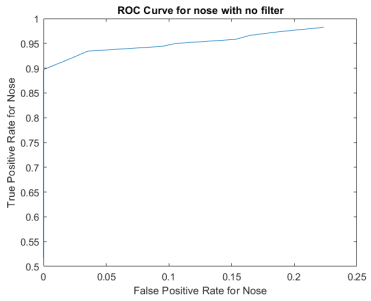

(a) ROC curve for nose with no filter

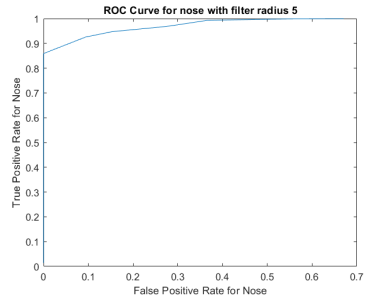

(b) ROC curve for nose with filter radius 5

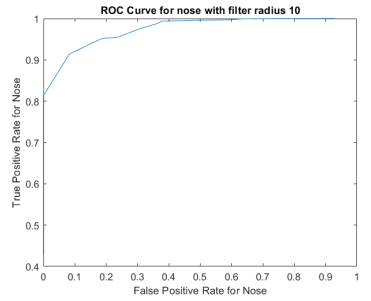

(c) ROC curve for nose with filter radius 10

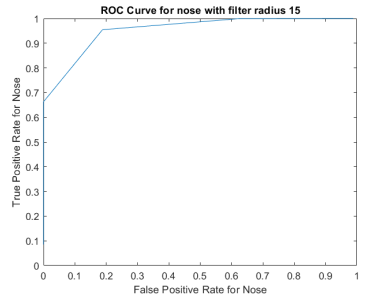

(d) ROC curve for nose with filter radius 15

Figure 21: ROC curves for the nose

B.5 ROC curves for the mouth

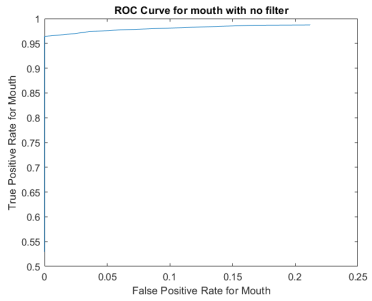

(a) ROC curve for mouth with no filter

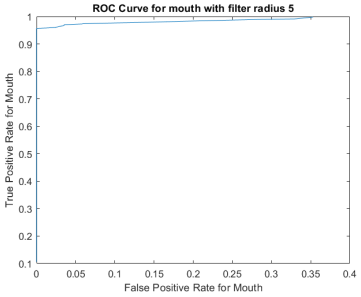

(b) ROC curve for mouth with filter radius 5

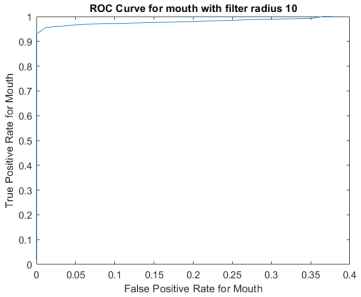

(c) ROC curve for mouth with filter radius 10

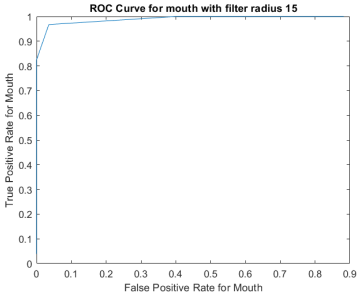

(d) ROC curve for mouth with filter radius 15

Figure 22: ROC curves for the mouth

B.6 ROC curves for the chin

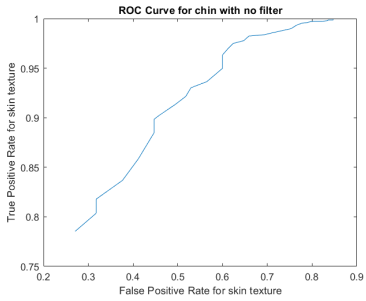

(a) ROC curve for chin with no filter

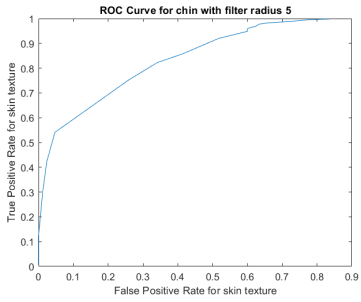

(b) ROC curve for chin with filter radius 5

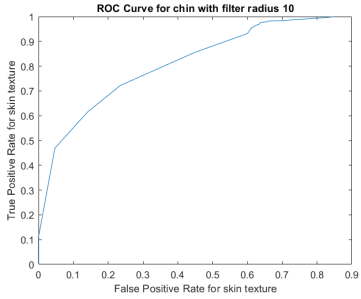

(c) ROC curve for chin with filter radius 10

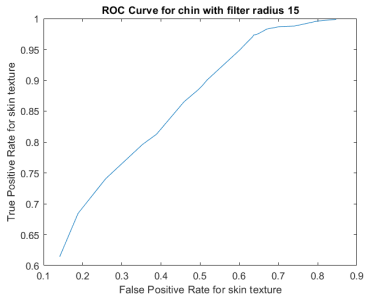

(d) ROC curve for chin with filter radius 15

Figure 23: ROC curves for the chin

C Histogram plots of scores for filtered DFTs

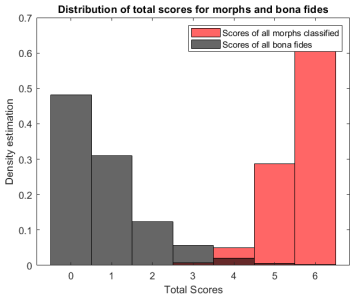

(a) Scores distribution for filter radius 5

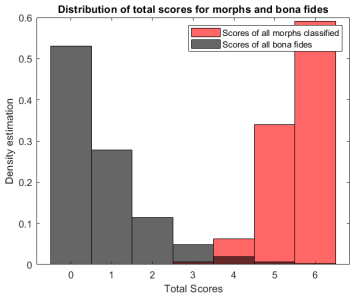

(b) Scores distribution for filter radius 10

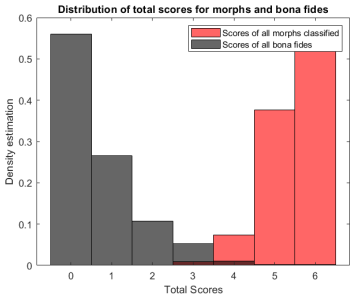

(c) Scores distribution for filter radius 15

Figure 24: positive decisions distribution for the filtered DFT of the eyebrows