

Detection of Morphed Face Images Using Discriminative Wavelet Sub-bands

Poorya Aghdaie, Baaria Chaudhary, Sobhan Soleymani, Jeremy Dawson, Nasser M. Nasrabadi
West Virginia University

Abstract—This work investigates the well-known problem of morphing attacks, which has drawn considerable attention in the biometrics community. Morphed images have exposed face recognition systems' susceptibility to false acceptance, resulting in dire consequences, especially for national security applications. To detect morphing attacks, we propose a method which is based on a discriminative 2D Discrete Wavelet Transform (2D-DWT). A discriminative wavelet sub-band can highlight inconsistencies between a real and a morphed image. We observe that there is a salient discrepancy between the entropy of a given sub-band in a bona fide image, and the same sub-band's entropy in a morphed sample. Considering this dissimilarity between these two entropy values, we find the Kullback-Leibler divergence between the two distributions, namely the entropy of the bona fide and the corresponding morphed images. The most discriminative wavelet sub-bands are those with the highest corresponding KL-divergence values. Accordingly, 22 sub-bands are selected as the most discriminative ones in terms of morph detection. We show that a Deep Neural Network (DNN) trained on the 22 discriminative sub-bands can detect morphed samples precisely. Most importantly, the effectiveness of our algorithm is validated through experiments on three datasets: VISAPP17, LMA, and MorGAN. We also performed an ablation study on the sub-band selection.

Index Terms—Morph detection, 2D discrete wavelet transform, information entropy, feature selection

I. INTRODUCTION

Morphing attack detection is of great significance in high-throughput border control applications. According to the CIA triad model, consisting of three main components, confidentiality, integrity, and availability of secure systems, morphed images violate the integrity of verification systems. A morphed image is generated using genuine face images from two different individuals. Because the resulting morphed image inherits characteristics of both subjects, it can be verified against both real subjects. Morphed images are generated using two approaches. In the first approach [1]–[3], two real face images are alpha blended in order to create a morphed image. To eliminate the ghosting effects in the morphed image, the average of the landmarks in both real images is used as the resulting landmark of the morphed image. In the second approach introduced in [4], a generative model, that is a Generative Adversarial Network (GAN), is trained to synthesize morphed images. Morph detection algorithms can be grouped into two main categories: single and differential morph detection. In the first category, an image under investigation is labeled as morphed or bona fide image, which is known as single image morph detection. In differential morph detection, a subject's image is

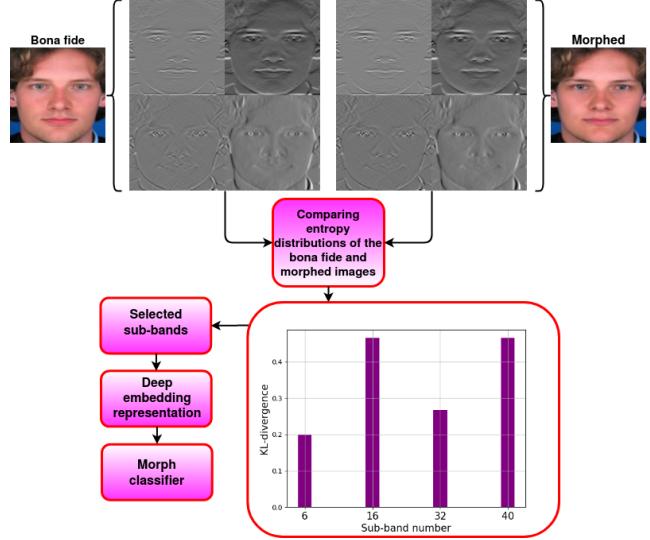


Fig. 1. A bona fide and a morphed image along with the four corresponding wavelet sub-bands. Using all the bona fide and morphed images in the dataset, 48 pairs of entropy distributions are found for bona fide and morphed images. Given a sub-band, dissimilarity between the two entropy distributions represents how discriminative that sub-band is with respect to morph detection. In the figure, sub-bands 16 and 40 are more discriminative than 6 and 32. A deep classifier is trained using the selected informative sub-bands.

compared with a live capture of the subject, and information from both images is used to detect morphed counterfeits.

To detect morphed images, some of the previous research efforts employ hand-crafted features such as Binarized Statistical Image Features (BSIF) [5], Scale Invariant Feature Transform (SIFT) [6], Speeded Up Robust Features (SURF) [7], (Local Binary Patterns Histogram) LBPH [8], Fused Local Binary Pattern (FLBP), and Histogram of Gradianets (HOG). Recently, Deep Neural Networks (DNNs) have proved to be promising in detecting morphed images [9], [10]. Thus far, no wavelet-based morph detection algorithm has been proposed. In this work, we propose a single image morph detector which can distinguish between a bona fide and a morphed face image. To do so, we train a deep neural network with a small number of selected discriminative wavelet sub-bands that are chosen according to the following criterion: the relative entropy between the entropy distribution of real faces and morphed faces is found for each of the wavelet sub-bands. The higher the value of the relative entropy for a given sub-band, the more discriminative that sub-band is for the task of

classification. Fig. 1 depicts our morph detection mechanism. Please note that only four wavelet sub-bands of a bona fide and its corresponding morphed image are selected in the figure for representation purpose. However, we consider all images in a given dataset to find the histogram of entropy for each sub-band. Experiments on three datasets, i.e., VISAPP17 [1], MorGAN [4], and LMA [4] verifies the performance of our morph detector. Standard quantitative measures, set forth by ISO/IEC 30107-3 [11], are used to evaluate the effectiveness of our proposed method. The first measure is Attack Presentation Classification Error Rate (APCER), which is the percentage of morphed images that are classified as bona fide. The second measure is Bona Fide Presentation Classification Error Rate (BPCER), which represents the percentage of bona fide samples that are classified as morphed. If we label the morphed class as positive and the bona fide class as negative, APCER, and BPCER are equivalent to false negative rate and false positive rate, respectively. The contributions of this paper are as follows: the most discriminative wavelet sub-bands are selected based on the KL-divergence between the two entropy distributions of both real and morphed images for the wavelet sub-bands. A DNN is trained using the selected informative wavelet sub-bands to detect morphed images. Finally, an ablation study is performed to show the effectiveness of our sub-band selection scheme for tackling detecting morph attacks.

II. RELATED WORK

Some morph detection techniques use hand-crafted features for training a classifier to identify morphed samples [4], [12], [13]. In addition, deep embedding features extracted using off-the-shelf DNNs can be utilized for training a morph detector [4], [13], [14]. Damer et al. [4] employ GANs for generating morphed face images. An SVM is trained for the task of morph detection using the crafted LBPH, as well as the extracted deep embedding feature vectors. One of the research efforts [15] adopts Photo Response Non-Uniformity (PRNU) to distinguish between real and morphed images. FLBP in color channels of HSV and YCbCr color spaces are studied in [16] as a method for detecting morphed images. One of the most important aspects of the morphing attacks is carefully selecting two bona fide subjects's face images such that the morphing attack looks highly photo realistic. In [17], the morph detection is investigated when the morphed images are generated using three different pairing protocols: (1) two similar images for morphing, (2) two random images, and (3) two dissimilar images. As a holistic approach for morph detection, fusion of the above-mentioned algorithms can be considered. In [13], two SVMs are trained using two different textures descriptors: LBPH, and BSIF. Another SVM is trained with the HOG, and deep embedding features are used to train another SVM. To integrate all approaches, the resulting scores from all the detectors are fused. Another work [10] employs a denoised version of an image to find the residual noise of the image which can be utilized for identifying morphed samples. The paper aggregates several denoised versions of an image

in the wavelet domain. Disentanglement of appearance and landmark is another method proposed for differential morph detection [18]. Interestingly, reflection inconsistencies are also employed to detect morphing attacks [19].

III. OUR FRAMEWORK

We employ undecimated 2D wavelet decomposition to address morphing attacks. Shannon entropy and Kullback-Liebler divergence [20] are utilized to identify the optimal discriminative sub-bands. In particular, the Shannon entropy [21] is used to measure embedded information in each sub-band of the wavelet decomposition. Since most of the morphing pipeline artifacts lie in the high frequency spectrum, we do not consider the Low-Low (LL) sub-band of the first level of decomposition to be decomposed further. Instead, the Low-High (LH), High-Low (HL), and High-High (HH) sub-bands are decomposed. After 3-level uniform decomposition, 48 sub-bands are obtained, for all of which the Shannon entropy is computed, and the distribution of the entropy is obtained for both real and morphed images for the three training datasets. The Kullback-Leibler divergence (relative entropy) is calculated between the entropy distribution of real and morphed sub-bands for each of the 48 sub-bands, and these 48 relative entropy values are sorted from highest to lowest. A final subset composed of 22 optimal discriminative sub-bands are selected that are used to train a DNN to detect morphed samples. As for the DNN, we employ a pre-trained Inception Resnet v1 architecture as our binary classifier.

A. Sub-band Selection Based on KL Divergence of Entropy Distributions

The pivotal point here is to distinguish morphed samples by leveraging the most discriminative sub-bands. To do so, we find the histograms of entropy of all 48 sub-bands for both bona fide and morphed images in the three datasets. Accordingly, 96 distributions are estimated using the histograms from the 48 sub-bands of both the bona fide and morphed presentations. The term \hat{f}_{b_i} represents the estimated distribution for the i^{th} sub-band pertinent to the bona fide images, and similarly, \hat{f}_{m_i} represents the estimated distribution for the i^{th} sub-band pertinent to the morphed images. The dissimilarity of the two probability distribution functions, namely $(\hat{f}_{b_i}, \hat{f}_{m_i})$ are calculated for all 48 sub-bands. The KL-divergence is the metric we employ to assess the dissimilarity between the distributions.

In order to select the most discriminative sub-bands, the KL-divergence values of each dataset are first normalized by removing the mean. The values are normalized to enable comparison of the distributions across the three datasets. Then, the zero-meaned values are averaged over the three datasets for each sub-band. The higher the KL-divergence value for a single sub-band, the more informative and discriminative that sub-band is in terms of classification. By choosing the sub-bands that are based on the highest average KL-divergence values from all datasets instead of each dataset separately, we can find the sub-bands that are discriminative across the

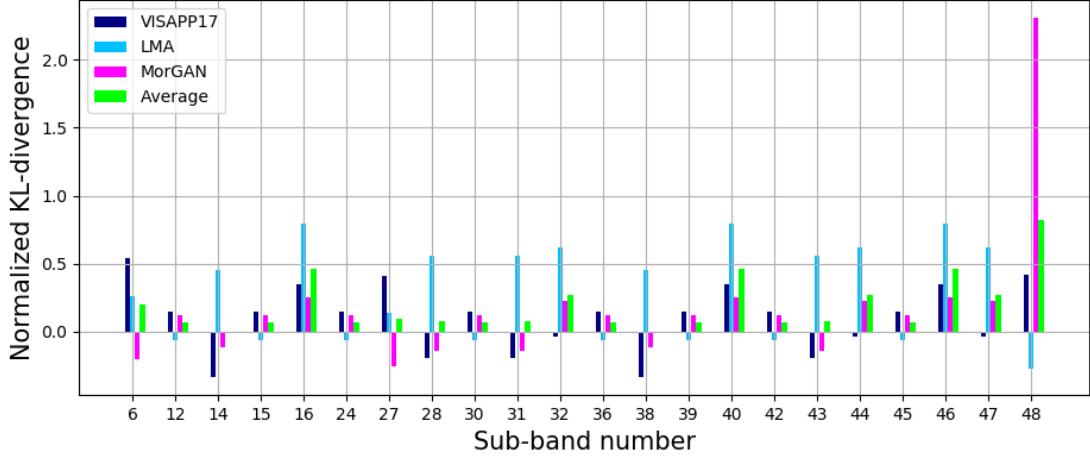


Fig. 2. Zero-meaned KL-divergence values in the top 22 most discriminative wavelet sub-bands for three datasets: VISAPP17, LMA, and MorGAN. The zero-meaned average of the KL-divergence values in each sub-band, as related to the three datasets, is represented in green.

Algorithm 1: Our Sub-band Selection

```

Input : Bona fide and Morphed Images
Output: A Set of Indices for Informative Sub-bands
 $\mathbb{I} = \{\}$ ; // index of sub-bands
for  $i = 1$  to  $48$ ; // sub-bands
  do
    for  $j = 1$  to  $3$ ; // datasets
      do
         $\hat{f}_{b_{ij}} \leftarrow \text{distribution}(H(S_{b_{ij}}))$ 
         $\hat{f}_{m_{ij}} \leftarrow \text{distribution}(H(S_{m_{ij}}))$ 
         $\mathbb{K}_{ij} \leftarrow D_{KL}(\hat{f}_{b_{ij}} \parallel \hat{f}_{m_{ij}})$ 
      end
    end
    for  $i = 1$  to  $48$  do
       $\bar{\mathbb{K}}_i \leftarrow \text{avg}(\mathbb{K}_{ij})$ 
      if  $\bar{\mathbb{K}}_i > \text{threshold}$  then
         $\mathbb{I} \leftarrow i$ 
      end
    end
  end

```

datasets, not just for a specific morphing technique. Fig. 2 shows the distribution of the zero-meaned KL-divergence values related to the 22 most discriminative wavelet sub-bands for the three morphed datasets, and their average values. Algorithm 1 illustrates our sub-band selection mechanism, in which $H(\cdot)$ represents the entropy function.

It is worth mentioning that the threshold for selecting the informative sub-bands is chosen using a data-driven method. After sorting the KL-divergence values from highest to lowest, different subsets of sub-bands are selected, e.g., top-5, top-10, and so forth. Suppose that the top-5 values from the set of informative sub-bands are selected. A DNN having input channel size of five is trained on the three training datasets combined, coined the *universal dataset*. The performance of the corresponding DNN is reported through Area Under

the Curve (AUC) metric using the validation portion of our universal dataset. Fig. 3 depicts the Area Under the Curve (AUC) when different numbers of sub-bands are chosen, based on which the optimal point for number of sub-bands is chosen as 22. The performance of the VISAPP17 dataset is consistent, irrelevant of the number of sub-bands used. This is primarily due to the small size of the VISAPP17 dataset (only 314 images—183 morphed and 131 real), and our DNN easily fits to VISAPP17 dataset regardless of the number of the selected sub-bands.

IV. EXPERIMENTAL SETUP

A. Datasets

Datasets used in this work are the VISAPP17 [1], MorGAN [4], and LMA [4]. The VISAPP17 dataset has been created using a landmark-based morphing attack, following by splicing, in which corresponding landmarks in two bona fide subjects are detected and the mean of each pair of the landmarks is calculated. Landmarks of each subject are then warped into the averaged landmark position, and the morphed image is generated using the blending of the two subjects' samples using triangulation [22] and then spliced into one of the contributing images. This technique aims to avoid artifacts that commonly arise from landmark manipulation, such as those that occur around the hairline. The MorGAN dataset is generated using a GAN. The encoder in a GAN can transform images to a latent space, and when two latent spaces related to two different subjects are combined, a morphed subject is synthesized. The LMA dataset is also generated using the landmark manipulation in two subjects' face images.

B. Training Setup

In this work, the Inception-ResNet-v1 architecture [23] is adopted as our DNN, which integrates the residual skips

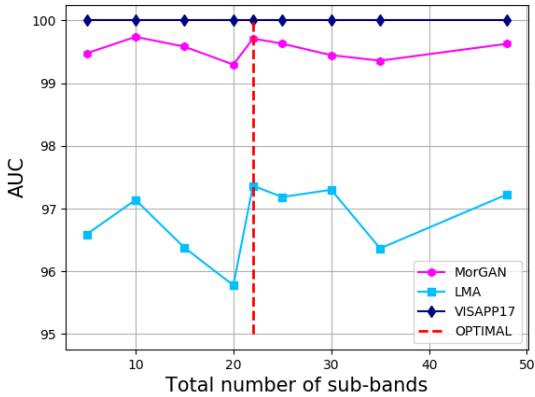


Fig. 3. Area under the curve versus number of sub-bands used in the training. Results indicate that 22 is the optimal number of sub-bands. After selecting 22 sub-bands, the performance does not increase significantly enough to validate using more sub-bands.

introduced in [24], and a revised version of Inception architecture [25]. We fine-tune an Inception-ResNet-v1, already pretrained on VGGFace2 [26]. The DNN is additionally fine-tuned with the obtained 22 discriminative wavelet sub-bands of the VISAPP17, MorGAN, and LMA datasets. An Adam optimizer [27] is employed for updating parameters of our network, and two 12 GB TITAN X (Pascal) GPUs accelerate our training.

C. Training/Testing Using Selected Sub-bands

In order to find the optimal number of sub-bands, we combine the three morph image datasets into a *universal dataset*. From this universal dataset, the training set consists of 1631 bona fide, and 1183 morphed samples. The validation set consists of 462 bona fide, and 167 morphed subjects. Moreover, the test set includes 1631 bona fide, and 1183 morphed images. We train several Inception-ResNet-v1 networks using the training portion of the universal dataset for a different number of chosen wavelet sub-bands. We assess the performance of the trained networks using the validation portion of the universal dataset. In other words, we do a search over the number of wavelet sub-bands, which is the input channel size of our convolutional neural network. Please note that the wavelet sub-bands are already sorted based on the corresponding KL-divergence values from highest to lowest. According to the sub-band selection scheme mentioned in section 3.B. and Fig. 3, which shows the performance of the trained classifier using different number of wavelet sub-bands, the optimal number of informative sub-bands is 22; thus, our final DNN has 22 input channels consisting of the top 22 most discriminative sub-bands.

The performance of our morph detector, and the baseline methods for comparison are summarized in the Table 1. Please note that we have considered all the possible training/testing scenarios using the three datasets, i.e., the VISAPP17, LMA, and MorGAN. The corresponding Detection Error Trade-off

Train	Test	Algorithm	D-EER	5%	10%
VISAPP17	VISAPP17	BSIF+SVM [5]	16.51	35.61	26.79
		SIFT+SVM [6]	38.59	82.40	75.60
		LBP+SVM [8]	38.00	77.10	67.90
		SURF+SVM [7]	30.45	84.70	69.40
		Ours	0.00	0.00	0.00
LMA	LMA	BSIF+SVM [5]	54.00	93.31	88.95
		SIFT+SVM [6]	37.00	79.00	70.00
		LBP+SVM [8]	33.00	71.80	59.90
		SURF+SVM [7]	39.30	86.10	75.70
		Ours	31.86	83.80	71.21
MorGAN	MorGAN	BSIF+SVM [5]	54.80	92.32	88.87
		SIFT+SVM [6]	58.00	96.10	89.90
		LBP+SVM [8]	40.00	76.90	67.40
		SURF+SVM [7]	40.30	83.00	74.00
		Ours	41.00	93.60	85.00
VISAPP17	VISAPP17	BSIF+SVM [5]	51.19	83.65	75.00
		SIFT+SVM [6]	38.00	90.80	86.30
		LBP+SVM [8]	36.60	77.80	71.80
		SURF+SVM [7]	30.80	70.00	65.60
		Ours	68.80	100.00	98.90
LMA	LMA	BSIF+SVM [5]	33.05	78.34	62.86
		SIFT+SVM [6]	33.30	83.40	72.00
		LBP+SVM [8]	28.00	58.60	51.40
		SURF+SVM [7]	37.40	79.50	70.00
		Ours	8.80	14.90	7.91
MorGAN	MorGAN	BSIF+SVM [5]	42.01	89.77	79.19
		SIFT+SVM [6]	50.70	95.00	89.80
		LBP+SVM [8]	35.00	72.60	61.30
		SURF+SVM [7]	41.27	84.60	78.00
		Ours	32.22	76.22	62.50
VISAPP17	VISAPP17	BSIF+SVM [5]	63.00	100.00	100.00
		SIFT+SVM [6]	42.00	92.40	84.00
		LBP+SVM [8]	42.32	84.70	79.30
		SURF+SVM [7]	31.40	74.00	55.70
		Ours	2.20	0.59	0.00
LMA	LMA	BSIF+SVM [5]	53.00	95.25	92.46
		SIFT+SVM [6]	40.20	90.70	80.00
		LBP+SVM [8]	39.18	75.90	67.7
		SURF+SVM [7]	39.40	81.00	71.60
		Ours	39.11	89.55	80.25
MorGAN	MorGAN	BSIF+SVM [5]	1.57	1.42	1.30
		SIFT+SVM [6]	43.50	93.20	84.20
		LBP+SVM [8]	20.10	52.70	32.30
		SURF+SVM [7]	39.95	80.00	72.60
		Ours	0.00	0.00	0.00

TABLE I
PERFORMANCE OF SINGLE MORPH DETECTION: D-EER%,
BPCER@APCER=5%, AND BPCER@APCER=10%.

(DET) curves are displayed in Fig. 4. In addition, we have trained the morph classifier using the training portion of the universal dataset, and the performance of that network is also evaluated using the testing portion of each individual dataset, as well as the universal dataset. The results of the training using the universal dataset, and the corresponding baseline methods are provided in the Table 2. Moreover, the related DET curves are shown in Fig. 5.

D-EER represents Detection Equal Error Rate, where APCER equals BPCER. BPCER5 designates BPCER rate for APCER=5%, and BPCER10 designates BPCER rate for

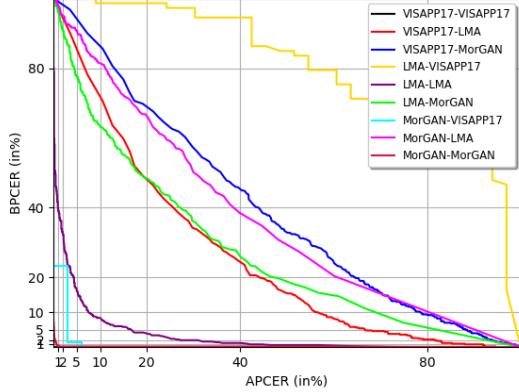


Fig. 4. DET curves when our morph detector is trained and tested on the selected 22-sub-band datasets. The legend represents train-test datasets.

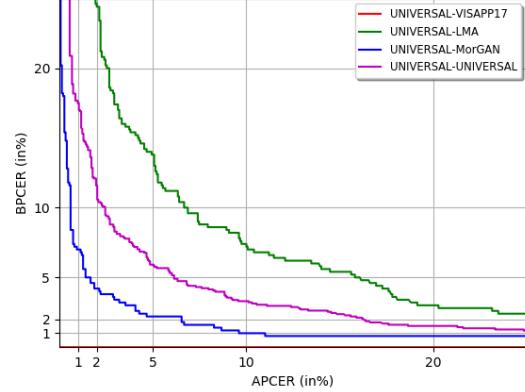


Fig. 5. DET curves when our morph detector is trained on the selected 22-sub-band universal datasets.

Train	Test	Algorithm	D-EER	5%	10%
VISAPP17	BSIF+SVM [5]	35.00	67.20	59.00	
	SIFT+SVM [6]	27.00	83.20	70.90	
	LBP+SVM [8]	37.67	72.50	59.50	
	SURF+SVM [7]	31.00	79.40	70.10	
	Ours	0.00	0.00	0.00	
LMA	BSIF+SVM [5]	30.00	70.42	57.60	
	SIFT+SVM [6]	28.31	67.70	50.00	
	LBP+SVM [8]	29.00	61.50	51.20	
	SURF+SVM [7]	33.40	74.50	62.70	
	Ours	8.61	12.93	7.05	
MorGAN	BSIF+SVM [5]	28.80	62.42	45.70	
	SIFT+SVM [6]	47.60	92.30	88.60	
	LBP+SVM [8]	31.20	62.00	55.60	
	SURF+SVM [7]	38.67	76.00	70.00	
	Ours	3.10	2.04	3.89	
Universal	BSIF+SVM [5]	23.74	51.42	38.67	
	SIFT+SVM [6]	37.21	87.45	76.71	
	LBP+SVM [8]	38.80	91.36	83.40	
	SURF+SVM [7]	36.00	75.50	65.76	
	Ours	5.45	5.70	3.19	

TABLE II

PERFORMANCE OF SINGLE MORPH DETECTION: D-EER%, BPCER@APCER=5%, AND BPCER@APCER=10%.

APCER=10%. A close scrutiny of the DET curves in Fig. 4 reveals that our morph detector can accurately detect morphed samples in both the VISAPP17, and MorGAN datasets when both training and testing data originate from the same dataset. Fig 5. also shows that our morph detector is able to detect most of the morphed samples in the VISAPP17, and MorGAN datasets when the classifier is trained on the training portion of the universal dataset.

D. Class Activation Maps

Class activation maps, set forth in [28], show the extent to which different regions in a given image contribute to the final classification decision for every class in an already trained DNN. After training an Inception-ResNet-v1 morph detector, class activation maps were constructed using the



Fig. 6. Class activation maps. Left: bona fide subject 1, middle: morphed subject, right: bona fide subject 2.

feature embeddings from the last layer before fully connected and softmax layers. The results are interestingly indicative of the likelihood that an image will be classified as morphed or bona fide. For example, in Fig. 6, the middle image, representing a morphed one, has many more affected areas than the other two bona fide images. This is an indicator that the middle image is the most likely image among the three images to be classified as morphed. Given that, our trained DNN using 22 discriminative wavelet sub-bands is effectively distinguishing morphed images from the non-morphed images.

V. ABLATION STUDY

In this section, the effect of sub-band selection is examined. To prove the effectiveness of band selection, a visualization method, namely t-SNE [29], is adopted. A total of 200 morphed, and 200 bona fide images are selected from the test set of MorGAN dataset. Fig. 7 shows the t-SNE visualizations for three scenarios using the MorGAN dataset, the first of which visualizes the original images, which is shown in the leftmost column. In the middle column, the 48 selected sub-band data is plotted. Finally, the 22 selected sub-band data is shown in the rightmost column. It is evident in Fig. 7 that sub-band selection contributes considerably to concentrating the morphed and bona fide data into separable clusters, which is highly desirable in terms of detecting morphed imagery.

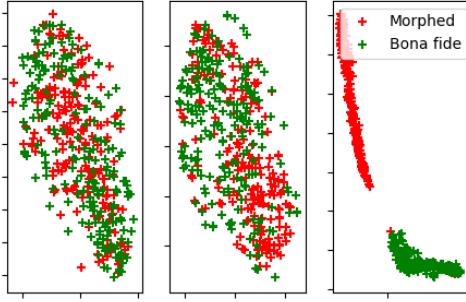


Fig. 7. T-SNE visualization for the original images (left), 48 sub-band data (middle), and 22 sub-band data (right). The 22 sub-bands evidently separate the morph and bona fide classes into very distinct clusters.

VI. CONCLUSION

In this paper, we proposed a framework to detect morphed face images using undecimated 2D-DWT. To select the optimal and informative bands, we found the distribution of the entropy for all the 48 wavelet sub-bands considering both the bona fide, and morphed images. The KL-divergence between the given distributions, integrated in a data-driven approach, led us to select the 22 most discriminative sub-bands. Furthermore, a close look at the presented results in Tables 1 & 2 highlights the fact that our morph classifier can identify morphed samples with a high accuracy in both the VISAPP17, and MorGAN datasets. Moreover, the ablation study on the sub-band selection substantiates the effectiveness of our method and shows that our trained DNN can map data samples to a new space where two bona fide and morphed classes are aggregated into two well-separated clusters.

REFERENCES

- [1] A. Makrushin, T. Neubert, and J. Dittmann, “Automatic generation and detection of visually faultless facial morphs.” in *VISIGRAPP (6: VISAPP)*, 2017, pp. 39–50.
- [2] C. Seibold, W. Samek, A. Hilsmann, and P. Eisert, “Accurate and robust neural networks for security related applications exemplified by face morphing attacks,” *arXiv preprint arXiv:1806.04265*, 2018.
- [3] ——, “Detection of face morphing attacks by deep learning,” in *International Workshop on Digital Watermarking*. Springer, 2017, pp. 107–120.
- [4] N. Damer, A. M. Saladié, A. Braun, and A. Kuijper, “Morgan: Recognition vulnerability and attack detectability of face morphing attacks created by generative adversarial network,” in *2018 IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS)*. IEEE, 2018, pp. 1–10.
- [5] J. Kannala and E. Rahtu, “Bsif: Binarized statistical image features,” in *Proceedings of the 21st international conference on pattern recognition (ICPR2012)*. IEEE, 2012, pp. 1363–1366.
- [6] D. G. Lowe, “Object recognition from local scale-invariant features,” in *Proceedings of the seventh IEEE international conference on computer vision*, vol. 2. Ieee, 1999, pp. 1150–1157.
- [7] H. Bay, T. Tuytelaars, and L. Van Gool, “Surf: Speeded up robust features,” in *European conference on computer vision*. Springer, 2006, pp. 404–417.
- [8] T. Ojala, M. Pietikainen, and D. Harwood, “Performance evaluation of texture measures with classification based on kullback discrimination of distributions,” in *Proceedings of 12th International Conference on Pattern Recognition*, vol. 1. IEEE, 1994, pp. 582–585.
- [9] L. Debiasi, N. Damer, A. M. Saladié, C. Rathgeb, U. Scherhag, C. Busch, F. Kirchbuchner, and A. Uhl, “On the detection of gan-based face morphs using established morph detectors,” in *International Conference on Image Analysis and Processing*. Springer, 2019, pp. 345–356.
- [10] S. Venkatesh, R. Ramachandra, K. Raja, L. Spreeuwiers, R. Veldhuis, and C. Busch, “Detecting morphed face attacks using residual noise from deep multi-scale context aggregation network,” in *The IEEE Winter Conference on Applications of Computer Vision*, 2020, pp. 280–289.
- [11] “International organization for standardization. iso/iec dis 30107-3:2016: Information technology biometric presentation attack detection part 3: Testing and reporting.” *Standard*, 2017.
- [12] U. Scherhag, C. Rathgeb, and C. Busch, “Towards detection of morphed face images in electronic travel documents,” in *2018 13th IAPR International Workshop on Document Analysis Systems (DAS)*. IEEE, 2018, pp. 187–192.
- [13] ——, “Morph detection from single face image: A multi-algorithm fusion approach,” in *Proceedings of the 2018 2nd International Conference on Biometric Engineering and Applications*, 2018, pp. 6–12.
- [14] R. Raghavendra, K. B. Raja, S. Venkatesh, and C. Busch, “Transferable deep-cnn features for detecting digital and print-scanned morphed face images,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE, 2017, pp. 1822–1830.
- [15] U. Scherhag, L. Debiasi, C. Rathgeb, C. Busch, and A. Uhl, “Detection of face morphing attacks based on prnu analysis,” *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 1, no. 4, pp. 302–317, 2019.
- [16] R. Raghavendra, K. Raja, S. Venkatesh, and C. Busch, “Face morphing versus face averaging: Vulnerability and detection,” in *2017 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2017, pp. 555–563.
- [17] N. Damer, A. M. Saladié, S. Zienert, Y. Wainakh, P. Terhöst, F. Kirchbuchner, and A. Kuijper, “To detect or not to detect: The right faces to morph,” in *2019 International Conference on Biometrics (ICB)*. IEEE, 2019, pp. 1–8.
- [18] S. Soleymani, A. Dabouei, F. Taherkhani, J. Dawson, and N. M. Nasrabadi, “Mutual information maximization on disentangled representations for differential morph detection,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021, pp. 1731–1741.
- [19] C. Seibold, A. Hilsmann, and P. Eisert, “Reflection analysis for face morphing attack detection,” in *2018 26th European Signal Processing Conference (EUSIPCO)*. IEEE, 2018, pp. 1022–1026.
- [20] S. Kullback and R. A. Leibler, “On information and sufficiency,” *The annals of mathematical statistics*, vol. 22, no. 1, pp. 79–86, 1951.
- [21] C. E. Shannon, “A mathematical theory of communication,” *Bell system technical journal*, vol. 27, no. 3, pp. 379–423, 1948.
- [22] D.-T. Lee and B. J. Schachter, “Two algorithms for constructing a delaunay triangulation,” *International Journal of Computer & Information Sciences*, vol. 9, no. 3, pp. 219–242, 1980.
- [23] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, “Inception-v4, inception-resnet and the impact of residual connections on learning,” in *Thirty-first AAAI conference on artificial intelligence*, 2017.
- [24] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [25] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.
- [26] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, “Vggface2: A dataset for recognising faces across pose and age,” in *International Conference on Automatic Face and Gesture Recognition*, 2018.
- [27] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [28] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, “Learning deep features for discriminative localization,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2921–2929.
- [29] L. v. d. Maaten and G. Hinton, “Visualizing data using t-sne,” *Journal of machine learning research*, vol. 9, no. Nov, pp. 2579–2605, 2008.