Face Recognition Systems: Are you sure they only consider your face?

Pavan Srihari Darbha, Mauro Conti, Eleonora Losiouk and Rajib Ranjan Maiti

***Abstract*—Face recognition has been one of the major biometric authentication procedures in smart devices that allows users to provide an additional layer of security for accessing their device. The accuracy of image similarity should depend on the face and its expression, as could be extracted from the whole image. Importantly, the background may have a substantial amount of additional information that can potentially pose a threat to the privacy of the user. In this paper, we report the impact of background on the recommended measure of similarity, Euclidean-L2, across different pictures that represent distinguishable emotions and image background. Additionally, we report that this impact of the background varies for different ethnic groups. Our findings are despite the fact that background should not matter for Face Recognition. For each facial image, we perform two preprocessings, gray-scaling and background whitening, and compute the similarity between the original image and the preprocessed image by using the DeepFace Face Recognition System. We have considered six data sets, i) containing 100 blurry images of one American man, ii) and iii) contained 100 images each of one American man in normal settings, iv) contained 50 each of East Asian men and women, v) contained 50 each of Indian men and women, and vi) contained 50 each of African or African-American men and women. We observe that gray scaling or background whitening images makes them dissimilar, often to the point of being unrecognisable. Overall, we report that the information contained in the background of a facial image can be significant and it can have different impacts across different skin complexions and facial structure. Importantly, our initial results bring up an important question of how to identify the images having a higher risk of exposing private information via the background of a facial image.**

***Index Terms*—Face Recognition, Mobile Biometric Authentication.**

I. Introduction

IMAGE Recognition is one of the most efficient ways for a computer to visually perceive objects in an environment. Face Recognition (FR) Systems consist of hardware and software components. The hardware components include cameras or lenses that take the picture or scan the face of the person, as well as a secure manner of storing data about the authorized personnel and a log of usage. Trusty OS with TEE (Trusted Execution Environment) [1] is one example of a secure OS and environment for FR or even Fingerprint Recognition. On the software side, programs are required to carry out the actual FR process and to accept queries and display results on the screen in use. FR Systems are used for automatic verification of users, usually taking digital photos or individual frames captured from video streams as input. Many government institutions, from law-enforcement to healthcare [2] and even private organizations use such a system for FR, often for identification of personnel with video cameras or in biometric systems to authenticate the users via cameras and scanners. Many prominent smartphone manufacturers, including Apple with iPhone X, Samsung with the Galaxy Note 9, and LG with G7 [3], also used FR as a means to allow end users to unlock their systems. FR Systems are also helpful in various other fields such as fighting crime and terrorism and enhancing security in real and virtual spaces. The system is expected to recognize where the face is in a photo and extract this area from the photo, compare it with existing faces in a pre-made database, and perform the verification test. This is, however, complex and problematic. People and sometimes even criminals often find methods to fool FR Systems using 3D images and faking depth in images [4]. There is also a huge variety of human faces depending on the area and conditions in which the photo is taken, including but not limited to camera performance, facial expressions, lighting, makeup, glasses, facial hair and more. Frequently, small variations in these conditions can affect the accuracy of the FR system [5].

Among the available biometric user authentication methods, FR and the more recent fingerprint unlock are popular in smartphones, tablets, and some laptops as well. FR has come under question in recent times during the pandemic as people questioned whether or not they could be recognized through masks [6]. The popular approach to FR was to use features of the human face and quantify them in a mathematical form. This involves graphs represented by vector maps and key points on the face treated as nodes. The human face was broken down into around 80 different parameters for features such as nose width, distance between eyes, height of eye sockets, bone structure of the face, width of the jaw, and many others for use as classification parameters. An image gallery is required to be created to serve as a set of models as a reference to the biometric matching process. Training Neural Networks to identify the facial area and to perform feature extraction is one common technique to allow for FR Systems to work.

To sum up, FR involves the following steps:

1) Capturing the image. The better quality of the image, the better, hence more powerful, cameras are preferred.
2) Face Detection. A complex process which deals with identifying a set of pixels in the image which contain

Pavan Srihari Darbha was an undergraduate student of Computer Science at BITS-Pilani, Hyderabad, India. e-mail: f20170011h@alumni.bits-pilani.ac.in

Mauro Conti (e-mail: conti@math.unipd.it) and Eleonora Losiouk (e-mail: elosiouk@math.unipd.it) are with the Department of Mathematics, University of Padua, Italy.

Rajib Ranjan Maiti is with the Department of Computer Science and Information Systems, BITS-Pilani, Hyderabad, India. e-mail: rajibrm@hyderabad.bits-pilani.ac.in

the face and which contain the background.
3) Feature Extraction. A mathematical representation of the biometric parameters used for verification is generated here.
4) Comparison of models. The mathematical representations of two images are compared VS one another by the use of some formulae for measuring *similarity* as a percentage or *distance* as a number.
5) Output Match or Mismatch. When the similarity is within a certain percentage or the distance is within a certain range or threshold value, the images are said to be of the same person and a match is declared. Human involvement is necessary to determine the threshold values in each case.

When using a smartphone camera for FR, the extra space that the camera captures can be a threat to privacy due to the fact that much of the background is also visible to the camera. Addressing this issue has involved the use of a rectangle on the screen to fit the user's face. Yet, privacy loss is still possible when the rectangle of the input face contains a lot of superfluous background space, as seen when the user holds the phone very far from their face.

To verify the impact of the background on FR, we whitened the background of images to see if the FR system would be affected by the change. We ran this test on 6 datasets, each containing 100 images. The first three datasets contained 100 images of one Caucasian man each. The last three were diverse sets, with 50 images each of men and women, totalling to 100 images of different people from different geographical locations and skin tones.

The distance, regardless of metric or dataset, was always 0 when an image was compared with itself. Upon applying the metrics, a disturbance was created, causing distance increases of around 0.3 on average for the Euclidean-L2 metric on Caucasians. The magnitude of this increase differed for different datasets, owing to existing issues with minority representation.

The results for the East Asians are similar to those of Caucasians, with a slightly higher average distance of 0.4. The results for Indians were worse still, with the distance exceeding 0.8 in the worst cases. For Africans/African-Americans, we saw that the fraction of images showing a distance greater than 0.5 is relatively higher. The greatest distance went beyond 1.0 as well.

Finally, prior studies also show that the performance of FR Systems is worse for underrepresented minorities due to Caucasians making up the majority of training data for the neural networks powering the FR systems [7]. We further explored this issue by applying filters like background-whitening and grayscaling and measuring the performance of FR Systems. We found out that, in the case of datasets representing minorities, the performance is worse than the one for datasets of Caucasians.

II. Background

FR Systems use machine learning. A neural network (NN) is trained on a predetermined set of data points. Here, images of people with varied facial angles and expressions, facial hair, and presence or absence of glasses, among other differences are used. The machine is given points when it correctly recognizes a person and points are subtracted when it fails to do so. It is then coded to maximise its points. A variety of points on the face are used for *Feature Extraction*, such as the tip of the nose, outline of the mouth and eyes, and so on. The key points are used to make a vector map of the subject's face. Then, a map is also created for the face to be tested against. Once the maps are prepared using the NN, they are compared through various distance metrics.

Distance metrics are used to calculate the difference or deviation between two vector maps using mathematical formulae. The distance output by the formula is then compared with a predetermined threshold distance. When the measured distance is below the threshold, the two images are said to be of the same person. It follows that exceeding the threshold causes the images not to be recognized.

Different metrics used by FR systems include *Cosine, Euclidean and Euclidean-L2*. Other popular metrics are *Mean Square Error, Manhattan, Correlation and Modified Manhattan*, as mentioned in [8].

Let two face images be denoted by F_a and F_b and called as source and test, respectively. Every image is converted to a fixed $n-$dimensional vector (n possibly being $224 \times 224 = 50176$) irrespective of its pixel dimension. Hence, $F_a = [f_1^a, f_2^a, ..., f_n^a]$ and $F_b = [f_1^b, f_2^b, ..., f_n^b]$. Then, one of the following similarity measures is considered on a pair of two n-dimensional vectors to compute the distance between the images.

Cosine(F_a, F_b). The computation can be done by performing a number of vector operations (such as dot product and magnitude), as follows.

$$a = \sum_{i=1}^{n} f_i^a \times f_i^b \quad b = \sum_{i=1}^{n} f_i^a \times f_i^a \quad c = \sum_{i=1}^{n} f_i^b \times f_i^b.$$

$$\text{Cosine}(F_a, F_b) = 1 - a/(\sqrt{b} \times \sqrt{c}).$$

The values of Cosine(F_a, F_b) vary in the range $\{0.0, 1.0\}$, where 0.0 indicates the most similar and 1.0 indicates the most dissimilar. The recommended threshold distance is 0.4.

Euclidean(F_a, F_b). We can subtract their components directly, followed by adding up the differences and squaring the result. This is the implementation of the Euclidean distance formula. The recommended threshold distance is 0.55.

$$e = \sum_{i=1}^{n} (f_i^a - f_i^b)^2 \qquad Euclidean(F_a, F_b) = \sqrt{e}.$$

Euclidean-L2(F_a, F_b). First, we normalize the vectors of F_a, F_b by turning them into unit vectors. Following this, the Euclidean Distance formula is applied to the normalized vectors. The recommended threshold distance is 0.75.

$$x = \sqrt{\sum_{i=1}^{n} f_i^a \times f_i^a} \qquad y = \sqrt{\sum_{i=1}^{n} f_i^b \times f_i^b}.$$

$$F_1^a = F_a/x \qquad F_1^b = F_b/y.$$

$$e = \sum_{i=1}^{n} (f_i^a - f_i^b)^2 \qquad E.L2(F_a, F_b) = \sqrt{e}.$$

Each of these metrics defines a threshold for identifying an image to be part of a class, i.e., a person in our case. We term this threshold as the identification threshold. Note that this threshold is expected to be the same across images of different persons. This is because an identified model for face identification is expected to be preinstalled on a mobile device and be personalized for a specific owner of it and the owner is not burdened to adjust the threshold.

III. Related Works

The paper by Y. Taigman *et al*. [9] is about Facebook's DeepFace system. It explains the system in detail and how it was the first system to reach or exceed near-human accuracy in identifying faces.

The work by Wang *et al*. [10] mentions the timeline of Deep FR in recent times and the success of both DeepFace and DeepID. It also illustrates how the focus on popular metrics has shifted over time. A very intriguing section of the paper talks about Privacy-preserving FR being an issue. An article by Gurovich *et al*. [11] is cited, which shows how even genetic information can be obtained from just an image, using Deep Learning.

The paper by Borade *et al*. [12] explores four popular metrics, of which Euclidean and Cosine were a part, and the influence of the metrics on the accuracy of the FR system.

The work by Varun *et al*. in [5] mentions that gender and ethnicity play an important role on perturbations in face. They posit that this could suggest bias in the training data.

Joy *et al*. [7] shows that existing FR Systems are biased toward dark-skinned people. Their work claims that the existing datasets under-represent minorities and women, leading to substantial bias in the learning system.

The results of Rajagopalan *et al*. [13] and Yoo *et al*. [14] show that preprocessing is applied to images before distance calculation to improve the accuracy. We show that, despite these results, there is an existing issue with DeepFace.

The result of Chen *et al*. [15] is that the background should, optimally, be ignored when using FR. Although this would be expected from an FR system, our observations on DeepFace do not align with this suggestion, especially the worsening of performance for minorities.

The poor performance of FR Systems for minorities was brought up by Klare *et al*. [16], stating that more inclusive training datasets would fix the issue. We build on this by showing that the performance further worsens when preprocessing is applied on the images of these minorities, though prior studies suggest that it should not.

IV. Datasets

We have used six data sets.

$\mathbf{D}_1$ contains 100 images of Robert Downey Jr., chosen judiciously out of a large pool of images in the PINS dataset [17], where every image contains the whole face of the person having various backgrounds. The images vary in face angle, hair orientation, presence or absence of glasses, and expression, and some images have additional faces while some others are blurry.

$\mathbf{D}_2$ contains 100 images of George W Bush selected carefully out of a pool of 530 images in the Labelled Faces in the Wild (LFW) dataset [18]. The images vary in light intensity and color, face angle, facial (uncommon) expressions, with or without glasses and having various backgrounds, such as more than one face, or containing the American national flag or the name of some conferences. Some images have unique attire, skiing clothes, for instance, along with protective goggles.

$\mathbf{D}_3$ contains 100 images of Donald Rumsfeld, selected by removing duplicate or similar images from a larger set in, again, LFW [18]. Though similar to $\mathbf{D}_2$, we chose this dataset to explore the impact of different types of dresses on Caucasian males.

The other three datasets are considered to find out if the issue worsens based on the person's ethnicity. We manually collected 100 images (50 men and 50 women) of three different ethnicities - Indian, East Asian (including Chinese, Korean, and Japanese), and African (including African-American).

$\mathbf{D}_4$ contains 100 images of African or African-American people, which includes 51 images of Nelson Mandela and the remaining include well-known women (Serena and Venus Williams and Michelle Obama), among others.

$\mathbf{D}_5$ contains 100 images of East Asian people, which includes 50 images of three well-known men (Jackie Chan, Wang Hao, and Ding Liren) and the remaining are of famous women including K-Pop stars, Japanese voice actresses and Chinese models.

$\mathbf{D}_6$ contains 100 images of Indian people, which contains 50 images of popular men's cricket players such as M.S. Dhoni and S.R. Tendulkar, and actors like Rajnikanth. The remaining 50 images are a mix of women's chess players like Harika Dronavalli and sportswomen such as Smriti Mandhana.

V. System Design and Implementation

To test our hypothesis of nonessential details affecting the distance metrics between two images, we first selected a few images and applied various filters, such as *Grayscale, Background-Whitening, Background-Replacement and Background-Grayscaling*. We made use of OpenCV2 [19] for grayscaling images and *segmentation* of images in conjunction with Numpy [20] for isolating background pixels and setting their values to white. We then used glob to automate the process for the entire folder of images in the dataset, while ensuring the face was never affected in the output images.

We then used DeepFace, with its VGG-Face2 setting, to calculate the distance between each image and its filtered version. Grayscale and Background-Whitening were found to be the most impactful. Hence, we proceeded to verify the findings on larger datasets. We iterated through the entire dataset, comparing each image with its own grayscaled and background-whitened versions separately, for each of the three metrics (i.e., Cosine, Euclidean and Euclidean-L2). The output distance values were stored in an array and written to a CSV file. The values in the CSV files were used to plot the graphs below. We focus on the Euclidean-L2 graphs since it is the recommended metric and the results are clearest.

Fig. 1: Euclidean-L2 Distances of 100 grayscale and background-whitened images from D1: Robert Downey Jr

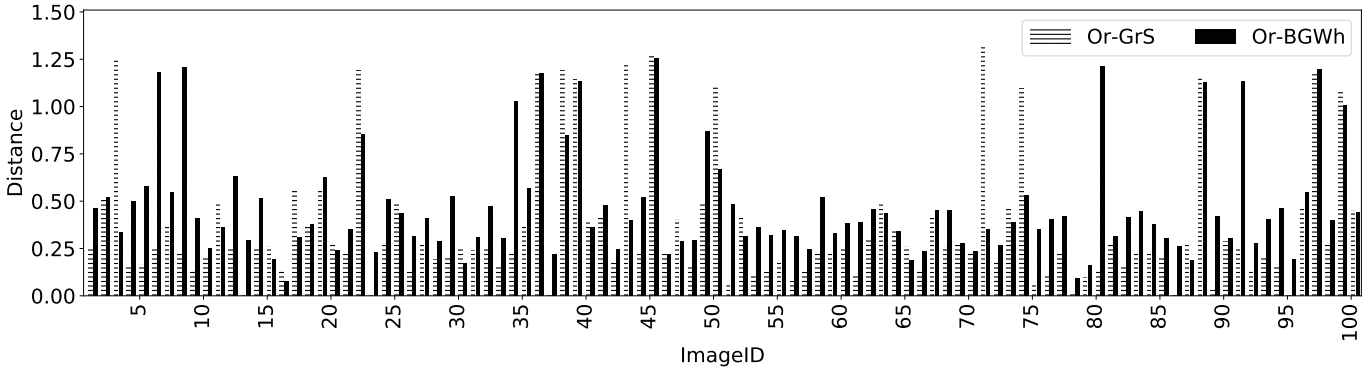

Fig. 2: Euclidean-L2 Distances of 100 grayscale and background-whitened images from D2: George Bush

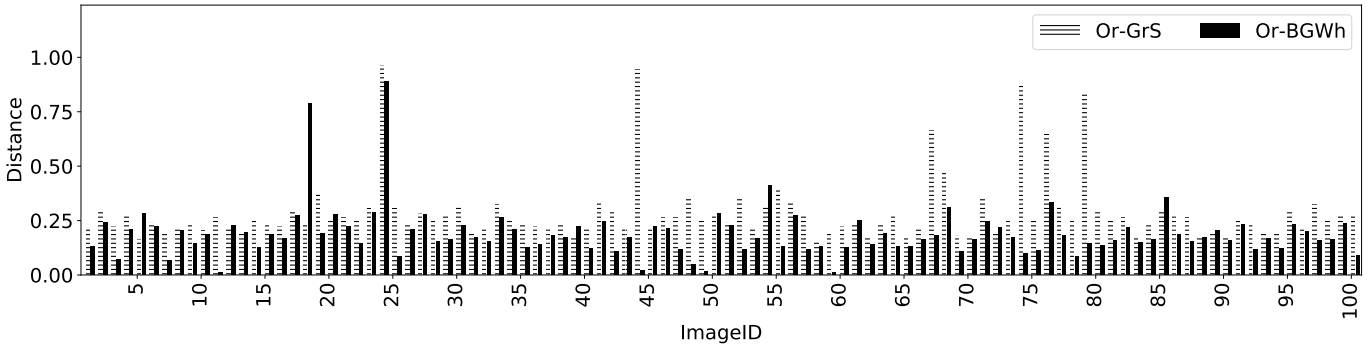

Fig. 3: Euclidean-L2 Distances of 100 grayscale and background-whitened images from D3: Donald Rumsfeld

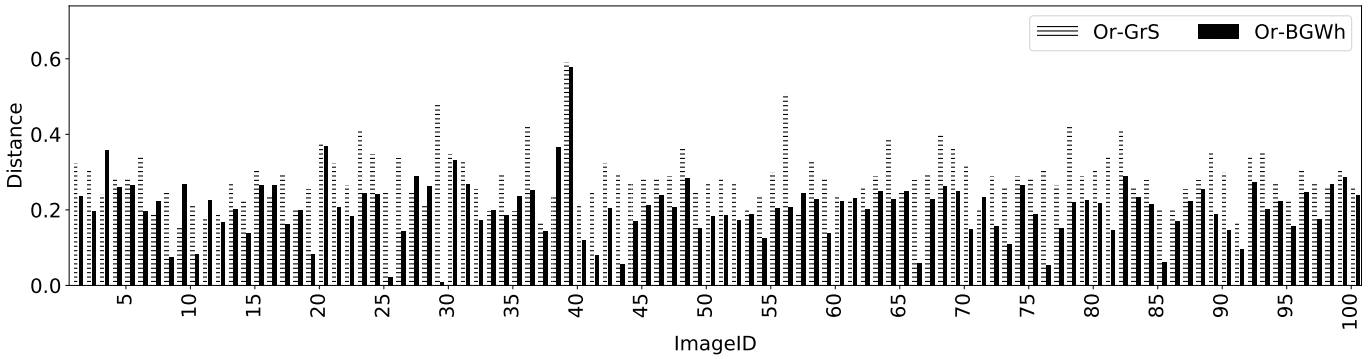

Fig. 4: Euclidean-L2 Distances of 100 grayscale and background-whitened images from D4: Africans

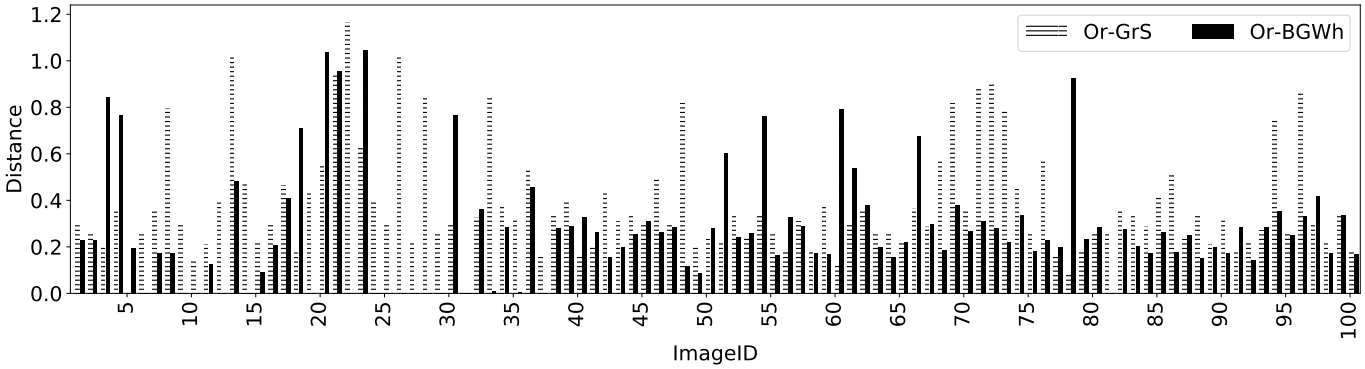

Fig. 5: Euclidean-L2 Distances of 100 grayscale and background-whitened images from D5: Asians

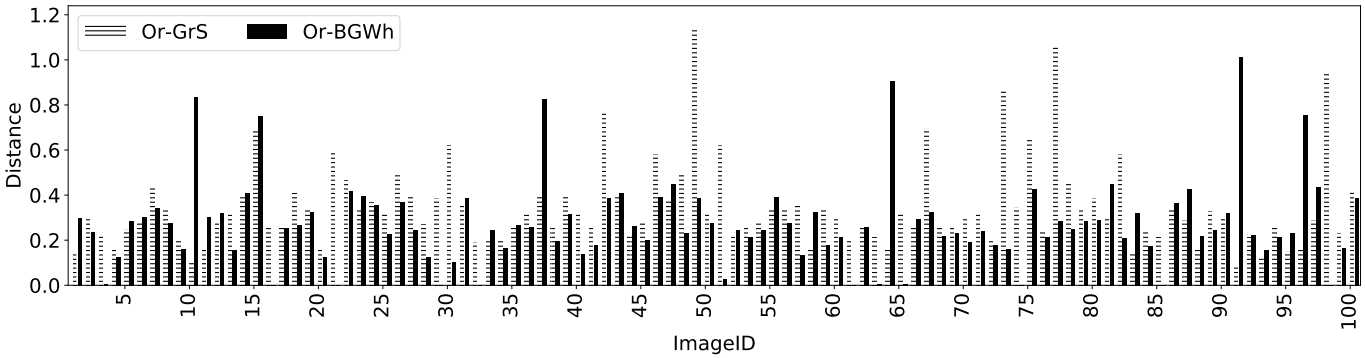

Fig. 6: Euclidean-L2 Distances of 100 grayscale and background-whitened images from D6: Indians

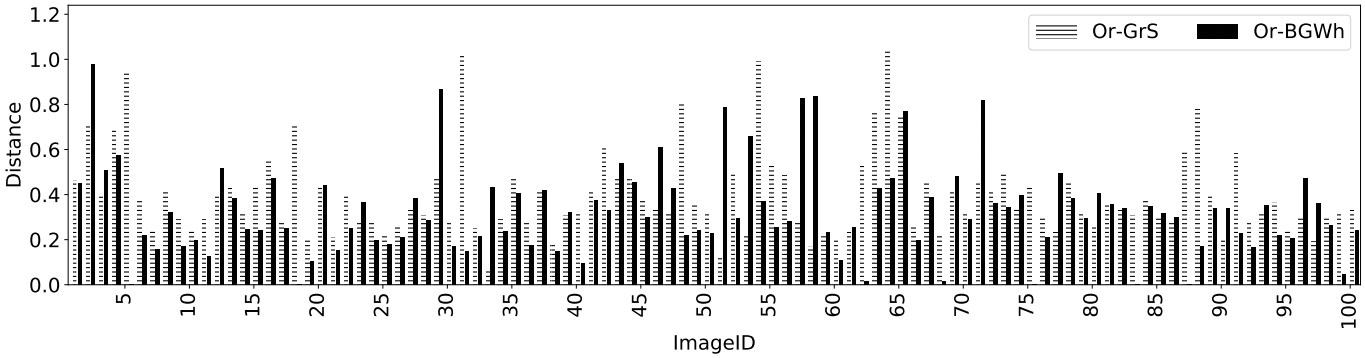

VI. Evaluation

We analyse the impact of the preprocessing steps, i.e., grayscaling and background whitening on the identification of a person based on facial images. Essentially, a facial image has a high possibility of containing additional information in the background of the image. A preprocessing step would try to eliminate any such information and focus on the portion of the image that contains faces. Thus, the impact of preprocessing on FR system implies the impact of such additional information in the background on the accuracy of FR. We consider two cases, (i) a similarity measure between a color image and its gray-scaled version and (ii) similarity between a color image and its background whitened version. Essentially, these two cases portray the impact of applying gray scaling and that of background whitening on the metrics used for image similarity. The images within a single data are given an arbitrary ordering of natural numbers for identification of a specific image in the data set.

The impact caused by the application of grayscale and background-whitening filters is most clear when the Euclidean-L2 metric is used. We first explore this impact on Caucasians, the primary ethnicity used for training NNs. In Figure 1, we show the EL2 results obtained by applying the filters on $\mathbf{D}_1$, which contains a wide variety of images of Robert Downey Jr. The average distance is around 0.5 for both filters. Before the filters were applied, the maximum distance was 0, which soared up to 1.3 in some cases after the filters were applied. Neither of the two filters seemed to have affected the system more than the other. Rather, both caused it serious issues with recognition. The main cause of these large distances is very likely to be the blurry and unorthodox nature of some of the images in this dataset.

Figure 2 shows the EL2 results obtained by applying the filters on $\mathbf{D}_2$, which contains images of George Bush. The average distance is around 0.25 for both filters. Before the filters were applied, the maximum distance was 0, which increased to almost 1.0 in some cases after the filters were applied. The grayscale filter caused worse performance than background-whitening for this dataset.

Figure 3 illustrates the EL2 results obtained by applying the filters on $\mathbf{D}_3$, which contains images of Donald Rumsfeld. The average distance is around 0.25 for both filters. Before the filters were applied, the maximum distance was 0, which increased to almost 0.6 in some cases after the filters were applied. For this dataset, the grayscale filter caused somewhat higher distance values than background-whitening almost throughout.

We next explore results on the performance of the FR system on minorities. Figure 4 shows the EL2 results obtained by applying the filters on $\mathbf{D}_4$, which contains images of men and women of a darker complexion. The average distance is above 0.3 for grayscaling and above 0.25 for background-whitening. Before the filters were applied, the maximum distance was 0, which increased to almost 1.3 in some cases after the filters were applied. Of the three datasets with minorities, the FR system had the worst performance for this set of images. Grayscaling appears to have had a more adverse effect on the performance of the system in general, for this dataset, in comparison to background-whitening.

Figure 5 illustrates the EL2 results obtained by applying the filters on $\mathbf{D}_5$, which contains images of men and women from Far East Asia. The average distance is around 0.3 for both filters. Before the filters were applied, the maximum distance was 0, which increased to almost 1.2 in some cases after the filters were applied. The trend was that grayscaling had a larger impact on the distance values, with some exceptions.

The graph in Figure 6 is based on the EL2 results obtained by applying the filters on $\mathbf{D}_6$, which contains men and women from India. The average distance is above 0.3 for both filters. Before the filters were applied, the maximum distance was 0, which increased to around 1.0 in some cases after the filters were applied. It appears that there is no trend between grayscale and background-whitening in terms of which filter had a greater impact From these results, both mathematically and visually, we show that the performance is indeed worse for minorities.

VII. Conclusion

FR has always been an interesting and potentially useful technology. However, in recent times, some issues have been detected. Our experiments make it clear that minor perturbations on images cause large changes in the distance as calculated by DeepFace, using the EL2 Distance Metric. Across the three out of six datasets which included 50 women each, we did not observe that the gender of the subjects could make any noticeable difference to the results, whereas, ethnicity turned out to be the major separator. There is a uniform trend in the performance of DeepFace based on ethnicity. The distances get progressively larger for less-represented minorities. This continues to hold after performing pre-processing, such as background-whitening and grayscaling. Based on our results, EL2 performs very poorly despite being the metric recommended, for use, by the creator. Our prior analysis indicated that the Cosine and Euclidean metrics are more reliable at preventing False Negatives, but EL2 is the most reliable at preventing False Positives. It is difficult to conclude on which metric is optimal for use when they all have their own efficacy, making it highly dependent on the use cases. Importantly, it is clear that the issues that exist with Neural Networks, for training FR Systems, get exacerbated when we consider under-represented ethnic groups. We speculate that this is the reason why we can observe numerous contemporary flagship phones lacking Face Unlock as a feature, while using Fingerprints instead. Until training datasets properly represent all ethnic groups, FR cannot be considered a foolproof, or even satisfactory, system.

REFERENCES

[1] ACRN, "Trusty tee and os," https://source.android.com/security/trusty.

[2] M. Gentzel, "Biased face recognition technology used by government: A problem for liberal democracy," Sep. 2021. [Online]. Available: https://doi.org/10.1007/s13347-021-00478-z

[3] cnet, "Flagships with face unlock," https://www.cnet.com/tech/mobile/10-best-phones-with-facial-recognition-iphone-x-note-9-galaxy-s9-lg-g7/.

[4] L. H. Newman, "Hackers trick facial-recognition logins with photos from facebook (what else?)," 2016. [Online]. Available: https://www.wired.com/2016/08/hackers-trick-facial-recognition-logins-photos-facebook-thanks-zuck/

[5] V. Chandrasekaran, C. Gao, B. Tang, K. Fawaz, S. Jha, and S. Banerjee, "Face-off: Adversarial face obfuscation," pp. 369–390, 2021.

[6] H. Walid, "Efficient masked face recognition method during the covid-19 pandemic," 07 2020.

[7] J. Buolamwini and T. Gebru, "Gender shades: Intersectional accuracy disparities in commercial gender classification," pp. 77–91, 23–24 Feb 2018. [Online]. Available: https://proceedings.mlr.press/v81/buolamwini18a.html

[8] H.-T. Nguyen, "Distance metrics for face recognition by 2d pca," 07 2015.

[9] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "Deepface: Closing the gap to human-level performance in face verification," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 1701–1708.

[10] M. Wang and W. Deng, "Deep face recognition: A survey," *Neurocomputing*, vol. 429, p. 215–244, Mar 2021. [Online]. Available: http://dx.doi.org/10.1016/j.neucom.2020.10.081

[11] Y. Gurovich, Y. Hanani, O. Bar, G. Nadav, N. Fleischer, D. Gelbman, L. Basel-Salmon, P. M. Krawitz, S. B. Kamphausen, M. Zenker, L. M. Bird, and K. W. Gripp, "Identifying facial phenotypes of genetic disorders using deep learning," *Nature Medicine*, vol. 25, no. 1, pp. 60–64, Jan. 2019. [Online]. Available: https://doi.org/10.1038/s41591-018-0279-0

[12] N. Borade, R. Deshmukh, and P. Shrishrimal, "Effect of distance measures on the performance of face recognition using principal component analysis," vol. 384, 09 2015.

[13] A. Rajagopalan, R. Chellappa, and N. Koterba, "Background learning for robust face recognition with pca in the presence of clutter," *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, vol. 14, pp. 832–43, 07 2005.

[14] S.-H. Yoo, S.-K. Oh, and W. Pedrycz, "Optimized face recognition algorithm using radial basis function neural networks and its practical applications," *Neural Networks*, vol. 69, 06 2015.

[15] L.-F. Chen, H.-y. Liao, J.-C. Lin, and C.-C. Han, "Why a statistics-based face recognition system should base its recognition on the pure face portion: A probabilistic decision-based proof," *Pattern Recognition*, vol. 34, pp. 1393–1403, 07 2001.

[16] B. F. Klare, M. J. Burge, J. C. Klontz, R. W. Vorder Bruegge, and A. K. Jain, "Face recognition performance: Role of demographic information," *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 6, pp. 1789–1801, 2012.

[17] Burak, "Pins dataset," 2020. [Online]. Available: https://www.kaggle.com/hereisburak/pins-face-recognition

[18] UMass, "Lfw dataset," 2008. [Online]. Available: http://vis-www.cs.umass.edu/lfw/

[19] G. Bradski, "The OpenCV Library," *Dr. Dobb's Journal of Software Tools*, 2000.

[20] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. Fernández del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, "Array programming with NumPy," *Nature*, vol. 585, p. 357–362, 2020.

Pavan Srihari Darbha graduated from BITS-Pilani, Hyderabad Campus, India, in June, 2021, with a Bachelor's degree in Computer Science. He was awarded the Aditya Birla Scholarship in 2017, which he retained throughout his undergrad years. He has been an intern under the guidance of Prof. Mauro Conti and Dr. Eleonora Losiouk at the University of Padua, Italy. His Thesis was 'The Privacy Issues of Mobile Face Recognition Systems' under the joint supervision of Prof. Conti and Dr. Rajib Ranjan Maiti, BITS-Pilani. He has also aided Prof. Conti's SPRITZ Group with their LOCARD Project. He intends to commence with his Master's in Computer Science in the Fall of 2022.

Mauro Conti is a Full Professor at the University of Padua, Italy. He is also affiliated with TU Delft and University of Washington, Seattle. He obtained his Ph.D. from Sapienza University of Rome, Italy, in 2009. After his Ph.D., he was a Post-Doc Researcher at Vrije Universiteit Amsterdam, The Netherlands. In 2011, he joined, as Assistant Professor, the University of Padua, where he became Associate Professor in 2015, and Full Professor in 2018. He has been Visiting Researcher at GMU, UCLA, UCI, TU Darmstadt, UF, and FIU. He has been awarded with a Marie Curie Fellowship (2012) by the European Commission, and with a Fellowship by the German DAAD (2013). His research is also funded by companies, including Cisco, Intel, and Huawei. His main research interest is in the area of Security and Privacy. In this area, he published more than 400 papers in topmost international peer-reviewed journals and conferences. He is Area Editor-in-Chief for IEEE Communications Surveys and Tutorials, and has been Associate Editor for several journals, including IEEE Communications Surveys and Tutorials, IEEE Transactions on Dependable and Secure Computing, IEEE Transactions on Information Forensics and Security, and IEEE Transactions on Network and Service Management. He was Program Chair for TRUST 2015, ICISS 2016, WiSec 2017, ACNS 2020, CANS 2021, and General Chair for SecureComm 2012, SACMAT 2013, NSS 2021 and ACNS 2022. He is Senior Member of the IEEE and ACM. He is a member of the Blockchain Expert Panel of the Italian Government. He is Fellow of the Young Academy of Europe. From 2020, he is Head of Studies of the Master Degree in Cybersecurity at University of Padua.

Eleonora Losiouk is an Assistant Professor at the University of Padua (Italy), working in the SPRITZ Group led by Prof. Mauro Conti. In 2018, she obtained her Ph.D. in Bioengineering and Bioinformatics from the University of Pavia (Italy). She has been a Visiting Fellow at EPFL in 2017. In 2020, she received the Seal of Excellence for her Marie Skłodowska-Curie individual project proposal and was awarded a Fulbright Fellowship for visiting ICSI, Berkeley (USA). Her main research interests regard the security and privacy evaluation of the Android Operating System.

Rajib Ranjan Maiti is currently an Assistant Professor in CSIS, BITS Pilani, Hyderabad campus in India. He has done his PhD in CSE at IIT Kharagpur, India. His research interest lies in the area of Cyber Security in IoT and CPS. He has published his research works in journals like Transaction on Mobile Computing, Computer Networks and Cybersecurity, and conferences like Esorics, WiSec and AsiaCCS. He is currently executing two sponsored projects related to cyber security, one funded by SERB, DST, India and the other funded by Axiado, India.