

FACIAL EXPRESSIONS AS A VULNERABILITY IN FACE RECOGNITION

A. Peña, A. Morales, I. Serna, J. Fierrez

Biometrics and Data Pattern Analytics Lab
Universidad Autonoma de Madrid, Spain

A. Lapedriza

Computer Science Dpt. / eHealth Center
Universitat Oberta de Catalunya, Spain

ABSTRACT

This work explores facial expression bias as a security vulnerability of face recognition systems. Despite the great performance achieved by state-of-the-art face recognition systems, the algorithms are still sensitive to a large range of covariates. We present a comprehensive analysis of how facial expression bias impacts the performance of face recognition technologies. Our study analyzes: i) facial expression biases in the most popular face recognition databases; and ii) the impact of facial expression in face recognition performances. Our experimental framework includes two face detectors, three face recognition models, and three different databases. Our results demonstrate a huge facial expression bias in the most widely used databases, as well as a related impact of face expression in the performance of state-of-the-art algorithms. This work opens the door to new research lines focused on mitigating the observed vulnerability.

Index Terms— Face Recognition, Expression

1. INTRODUCTION

The progress of face recognition techniques have allowed practical uses of these technologies, such as access control in private places (e.g. gyms or office buildings), safely unlocking your smartphone, or systems that facilitate automatic identity verification in the airports' border control. With the emergence of deep learning architectures, the accuracy of these technologies have largely surpassed human-level performance [1]. However, despite these impressive results, some studies show that deep learning based face recognition systems remain significantly sensitive to different sources of variation in face images, commonly referred as covariates [2, 3]. This has motivated researchers to study ways to reduce the impact of these factors for years.

On a different front, researchers have demonstrated that face recognition technologies are vulnerable to different types of attempts to intervene the automatic process of recognition, which have been traditionally studied as *attacks* [5]. These attacks can be divided into digital or physical attacks. Digital attacks are those intended to manipulate the data flow of

the system [6]. Physical attacks, also known as spoofing attacks, refer to face manipulations made before the data acquisition [7, 8]. The most common physical attack is the usage of tools (e.g. masks or makeup) to change our appearance [9, 10]. As shown in Fig. 1, the usage of accessories to disguise your real identity (e.g. scarf or glasses) directly affects the output of a face recognition system, a fact largely studied in forensic face recognition [11]. However, we can also observe another easy way to interfere the output of the Face Recognition system, whose impact is at least comparable to the usage of disguising tools: changing your facial expression. Although the influence of facial expressions on face recognition performance is well-known [12], they have typically been studied as another covariate associated to pose variations [13]. To the best of our knowledge, there are no works analyzing the impact of facial expression bias in current deep face recognition technology, and the vulnerability they may suppose to such systems.

In this work, we focus on the prototypical facial expressions of the 6 basic emotions: happiness, sadness, anger, surprise, disgust, and fear. In Sec. 2 we describe the different databases and algorithms used in our experiments. Then, we quantify the significant facial expression biases present in the most popular face recognition datasets (Sec. 3). To understand to what extent these facial expression biases affect the performance of face recognition models we performed a series of experiments in controlled scenarios. Our results (Sec. 4) demonstrate that popular face models trained solely for face identity recognition have self-learned emotional responses¹. We evaluate empirically the impact of such self-learned emotional responses in the security of face recognition systems using 3 public databases and 3 face recognition models.

2. DATA AND METHODS

In order to conduct a comprehensive study of the impact of facial expressions on face recognition systems, we carry out

¹An emotional response is defined here as an output of a pre-trained model that varies depending of the face expression of the input image [14].




Fig. 1. Matching score (Euclidean distance computed on ResNet-50 [4] embeddings) between a Reference image (grey dotted line) and 10 Test images including different facial expressions (orange line) and accessories to obfuscate the face (black line).

our experiments using different state-of-the-art face recognition pre-trained models:

- VGG16 [15]: This model has 5 convolutional blocks that use small convolutional filters in all layers to effectively increase the depth of the network. The model was trained with the VGGFace2 database [16] and the MTCNN face detector [17].
- ResNet-50 [4]: This CNN comprises 50 layers including residual connections. We use the model trained for face recognition with VGGFace2 [16] and the MTCNN face detector [17].
- LResNet100E-IR [18]: This model is a residual network [4] with 44M parameters, which use a more advanced residual unit setting [19]. The model was trained with ArcFace loss, the MS-Celeb-1M ArcFace dataset [18] and the RetinaFace detector [20].

All three models were used as feature extractors, computing the matching score between two face images as the Euclidean distance between its two feature vectors. We used three public databases to evaluate the effect of changes in facial expression in face recognition performance: CFEE [21], CK+ [22], and CelebA [23]. Note that, while CelebA is a large database collected using search engines, both CFEE and CK+ were collected under lab-controlled conditions, and therefore they are ideal to analyze the isolated effect of facial expression on the face recognition pipeline.

3. EXPRESSION BIAS IN FACE DATASETS

One of the main challenges for training a face recognition system based on Deep Models is the need for millions of images from thousands of users involving realistic acquisition conditions (e.g. changes in pose, illumination, or image resolution). This led to the design of large-scale databases using search engines and semi-automated processes to collect images from the Internet, movies or social media in an unconstrained way (e.g. VGGFace2 [16] or MS-Celeb-1M [24]). Although these databases have significantly contributed to

make great advances in the field, they may present biases that can impact the performance of the trained systems.

Inspired by the work of [28], where the authors analyze the demographic statistics of multiple face databases, we conducted a similar experiment to study the facial expression bias in popular face databases. For that, we used the COST Affective, whose reliability in the expression classification task has already been tested [29], to classify database images into 7 facial expressions (6 basic emotions plus neutral face). Our hypothesis is that the source of the images of these databases can be biased towards certain facial expressions, with some expressions significantly more frequent than others. Tab. 1 summarizes the facial expression distribution in each database. Note that there is a strong facial expression bias in all the evaluated databases. In all cases, the proportion of Neutral images exceeds 60%, reaching a maximum value of 83.7% in the MS-Celeb-1M database [24]. The second most common facial expression is Happy. In fact, for all the datasets, $\sim 90\%$ of the images show either Neutral or Happy.

As for the other 5 facial expressions, notice that Surprised and Disgusted rarely exceed 6%, while Sad, Fear, and Anger have in general a very low representation, often below 1%. This remarkable under-representation of some facial expressions produces the following drawbacks: i) on the one hand, models are trained using highly biased data that result in heterogeneous performances; ii) on the other hand, technology is evaluated only for mainstream expressions hiding its real performance for images with some specific facial expressions.

4. IMPACT OF FACIAL EXPRESSION IN FACE RECOGNITION SYSTEMS

Our hypothesis is that the facial expression biases of the training set can impact the robustness of the face recognition systems. Concretely, the biases revealed in Sec. 3 make us think that the models will work better for faces showing Neutral or Happy expressions, which are the most common expressions.

We start our analysis on CFEE [21], which contain images collected in a very controlled environment with uniform illu-

Database	#images	Neutral	Happy	Sad	Anger	Surprised	Disgusted	Fear
MS-Celeb-1M [24]	8.5M	83.7%	5.7%	0.2%	3.4%	2.2%	4.6%	~ 0.0%
MegaFace [25]	4.7M	82.0%	4.5%	0.1%	7.0%	1.3%	5.0%	~ 0.0%
CelebA [23]	203K	62.2%	33.3%	0.1%	0.5%	1.6%	0.9%	~ 0.0%
IJB-C [26]	21K	66.2%	26.9%	0.1%	0.6%	2.6%	2.0%	0.1%
LFW [27]	13K	61.2%	28.0%	0.3%	1.8%	3.1%	4.4%	0.0%
VGGFace2 [16]	3.3M	64.5%	28.2%	0.2%	0.4%	3.3%	2.0%	0.1%

Table 1. Facial expression labels of images in the most popular state-of-the-art face databases. The facial expressions were labeled using the Affectiva software (<https://www.affectiva.com/>).




Fig. 2. Genuine (continuous line) and impostor (dashed line) matching score distributions by facial expression on the CFEE database [21] for the face matchers: (Left) VGG16 [15], (Center) ResNet-50 [4], and (Right) LResNet100E-IR [18].

mination, controlled pose, and without occlusions. Each user in CFEE presents one image unequivocally representing each of the 22 expressions in the dataset. Thus, the impact of the facial expression can be isolated from the rest of covariates (e.g. pose, demographic attributes, illumination). We conducted an authentication experiment including seven facial expressions (i.e. Neutral expression and the six basic expressions defined in the FACS manual [14]), where we analyze both genuine and impostor scores distributions by comparing each facial expression with the rest. CFEE is composed of 230 users, all of them having an image representing each expression. Thus, all the distributions were computed using the same number of image pairs: 230 for each genuine distribution and 230×229 for each impostor distribution.

Fig. 2 presents the genuine and impostor matching scores distributions (pairs of faces), using Neutral expression as reference, computed with the three face recognition algorithms presented in Sec. 2. Notice that, while the genuine distributions do not include the case Neutral vs Neutral due to the limitations of CFEE (only one image per expression), the impostor distributions contemplate this scenario. We observe that, while the genuine distributions are clearly influenced by the facial expressions being compared, the impostor distributions barely change across expressions. In both ResNet-50 and LResNet100E-IR the impostor distributions are practically identical. The impostor distributions obtained with VGG16 show more overlap than the others, probably because the poorest performance of that model in comparison with ResNet-50 and LResNet100E-IR. At the same time, if we fo-

cus on the genuine distributions in Fig. 2, we clearly observe a difference among expressions in all three cases. For the rest of experiments we will focus on these genuine differences.

Tab. 2 presents the average genuine score for each distribution. Considering Neutral as the reference expression, the Happy, Sad, and Surprised distributions show the lowest scores, while the Disgusted and Anger distributions tend to higher values. For the three models in Fig. 2 we observe that all the distributions seem to have a similar deviation, except for the case of Happy, which shows a higher concentration of scores around the average. LResNet100E-IR is a state-of-the-art face recognition model with very competitive performances in a large list of databases and protocols. However, the results demonstrate that its output is largely affected by the facial expression, with genuine score variations of more than 30% across expressions in some cases. On the other hand, several distributions computed by VGG16 exceed the score threshold of 1, where the impostors begin, having a large overlap in the Anger and Disgust distributions.

Regarding the genuine score when non-neutral facial expressions are used as the reference (see Tab. 2), the observed trend is similar to the trend observed with the Neutral expression. The genuine distributions depend on the expressions compared in all three methods, with VGG16 having the greatest vulnerability to this effect (i.e. higher average score). In general, Anger, Sad, and Fear expressions work well together. Surprise expression has a better affinity with Fear and Neutral expressions, while Disgusted ones obtain the worst results with almost all expressions. Happy expressions seem to work

Met.	Av. Genuine Score Rank-1 in % (Ref. Expression = N)					
	H	S	A	Su	D	F
(a)	.7796.0	.7390.3	.8677.3	.7389.4	.9650.3	.7884.2
(b)	.50100	.52100	.6197.1	.50100	.6794.4	.5598.8
(c)	.55100	.53100	.59100	.5199.7	.6599.3	.56100
Met.	Av. Genuine Score Rank-1 in % (Ref. Expression = H)					
	N	S	A	Su	D	F
(a)	.7796.9	.8784.3	.9560.0	.8876.4	.9755.2	.8684.8
(b)	.50100	.6297.6	.6797.3	.6297.6	.6893.0	.6098.4
(c)	.55100	.6599.6	.6899.6	.6499.6	.6999.3	.63100
Met.	Av. Genuine Score Rank-1 in % (Ref. Expression = S)					
	H	N	A	Su	D	F
(a)	.8773.4	.7394.2	.7786.8	.8479.1	.9256.4	.7492.1
(b)	.6297.2	.5299.5	.5698.0	.6197.6	.6689.2	.5499.3
(c)	.6599.2	.54100	.5599.6	.6099.6	.6599.2	.56100
Met.	Av. Genuine Score Rank-1 in % (Ref. Expression = A)					
	H	S	N	Su	D	F
(a)	.9554.1	.7788.0	.8680.4	.9644.4	.8769.1	.9078.4
(b)	.6795.3	.5697.1	.6197.4	.7186.6	.6283.5	.6595.6
(c)	.68100	.55100	.59100	.6699.7	.6099.2	.63100
Met.	Av. Genuine Score Rank-1 in % (Ref. Expression = Su)					
	H	S	A	N	D	F
(a)	.8862.1	.8474.7	.9643.7	.7281.0	.9937.3	.7366.0
(b)	.6295.1	.6195.6	.7187.0	.5096.5	.7183.5	.5290.5
(c)	.6499.7	.60100	.6699.6	.5199.9	.7099.3	.53100
Met.	Av. Genuine Score Rank-1 in % (Ref. Expression = D)					
	H	S	A	Su	N	F
(a)	.9741.9	.9253.9	.8765.5	.9935.8	.9640.4	.9543.0
(b)	.6886.0	.6685.0	.6293.5	.7177.8	.6790.3	.6783.4
(c)	.6999.3	.6599.6	.6099.2	.7099.3	.6599.7	.67100
Met.	Av. Genuine Score Rank-1 in % (Ref. Expression = F)					
	H	S	A	Su	D	N
(a)	.8675.2	.7491.1	.9061.9	.7388.0	.9552.9	.7883.0
(b)	.6098.8	.5499.2	.6592.5	.5298.4	.6787.6	.5598.6
(c)	.6399.2	.5699.6	.6399.2	.5399.6	.6799.2	.5699.4

Table 2. Average genuine score in CFEE [21], and Rank-1 statistics (% in subscript) in CFEE, CK+ [22] and CelebA [23] with different face expressions as reference for the face matchers: (a) VGG16 [15], (b) ResNet-50 [4], and (c) LResNet100E-IR [18].

well with Neutral expressions (i.e. the two most common expressions in training databases, as analyzed in Sec. 3).

This impact in the genuine matching score distributions suppose a potential vulnerability, as it may influence the probabilities to be correctly identified by these systems. To study this possibility, we designed an identification experiment with the subjects in CFEE and CK+ [22], a total of 350 different identities and 16K different images. In order to increase the background set, we also add a set of 15K images from the CelebA [23]. We are aware that the characteristics of CelebA dataset are very different to those from CFEE or CK+. However, CelebA includes good quality images that can serve to increase the difficulty of the experiment. For each user in CFEE, we selected a specific expression image as query sample, computing the matching score with all the images belonging to the rest of users in the background set. We then compare the background scores with each of the user’s genuine scores, computed by comparing the query sample with

the other 6 basic facial expressions, and extract Rank-1 statistics individually for each facial expression.

Tab. 2 presents Rank-1 results by facial expression, using the three face recognition algorithms with their respective face detectors. We observe lower Rank-1 results when using VGG16 in all cases. Both ResNet-50 and LResNet100E-IR obtained good results, being slightly higher in the second one. The Rank-1 results follow the trend seen in the case of the average genuine scores, maintaining the relationships between expressions previously mentioned. The Disgusted expressions obtained the worst results in all cases, confirming the low affinity with the other expressions that we observed in the scores distributions. The drop of performance is clear for VGG-16 and ResNet-50 models, with Rank-1 differences between expressions up to 50%. These differences are moderate for the LResNet100E-IR model. The high quality standards of CFEE images are not a challenge to a state-of-the-art model. However, the drop of performance in LResNet100E-IR exists and varies between 0 – 1%, even in this simple environment. It is expected that this impact increases in a more challenging benchmark where expressions and other covariates affect the performance at the same time.

5. CONCLUSIONS

This work analyzes the impact of facial expressions bias on face recognition systems. Our main findings are: i) The most popular face recognition databases systematically present huge facial expression biases. Therefore, new public databases are needed to overcome the scarce representation of some facial expressions in existing databases. ii) Facial expression bias affects the performance of genuine comparisons with variations in scores of up to 40%. Facial expression bias does not affect impostor comparisons. Thus, we can discard impersonation attacks based on facial expression manipulation. However, obfuscation of identity based on manipulation the face expression is easy and particularly harmful.

As a result of this work, we strongly advocate for reducing the facial expression bias in future face recognition databases, and further development of bias-reduction methods applicable to existing databases and existing models already trained on biased datasets [30, 31, 32]. The application of face manipulation techniques [33, 34, 35] could serve to enhance the facial expression availability on existing database by introducing new synthetically generated expressions. Analyzing varying facial expressions in time is also in our plans [36].

6. ACKNOWLEDGMENTS

Supported by projects: TRESPASS-ETN (MSCA-ITN-2019-860813), PRIMA (MSCA-ITN-2019-860315), BIBECA (RTI2018-101248-B-I00 MINECO/FEDER), and RTI2018-095232-B-C22. A. Peña is supported by a research fellowship from MINECO.

7. REFERENCES

- [1] R. Ranjan et al., “Deep learning for understanding faces: Machines may be just as good, or better, than humans,” *IEEE Signal Processing Magazine*, vol. 35, no. 1, pp. 66–83, 2018.
- [2] E. Gonzalez-Sosa, J. Fierrez, et al., “Facial soft biometrics for recognition in the wild: Recent works, annotation and cots evaluation,” *IEEE Trans. on Information Forensics and Security*, vol. 13, no. 8, pp. 2001–2014, August 2018.
- [3] B. Lu, J. Chen, C. D. Castillo, and R. Chellappa, “An experimental evaluation of covariates effects on unconstrained face verification,” *IEEE Trans. on Biometrics, Behavior, and Identity Science*, vol. 1, no. 1, pp. 42–55, 2019.
- [4] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *IEEE CVPR*, 2016.
- [5] J. Galbally, J. Fierrez, and J. Ortega-Garcia, “Vulnerabilities in biometric systems: Attacks and recent advances in liveness detection,” in *Proc. Spanish Workshop on Biometrics*, 2007.
- [6] J. Galbally, C. McCool, J. Fierrez, et al., “On the vulnerability of face verification systems to hill-climbing attacks,” *Pattern Recognition*, vol. 43, no. 3, pp. 1027–1038, March 2010.
- [7] J. Galbally, S. Marcel, and J. Fierrez, “Biometric anti-spoofing methods: A survey in face recognition,” *IEEE Access*, vol. 2, pp. 1530–1552, 2014.
- [8] J. Galbally, S. Marcel, and J. Fierrez, “Image quality assessment for fake biometric detection: Application to iris, fingerprint and face recognition,” *IEEE Trans. on Image Processing*, vol. 23, no. 2, pp. 710–724, February 2014.
- [9] A. Hadid, N. Evans, et al., “Biometrics systems under spoofing attack: An evaluation methodology and lessons learned,” *IEEE Signal Processing Magazine*, vol. 32, no. 5, pp. 20–30, 2015.
- [10] J. Hernandez-Ortega, J. Fierrez, A. Morales, and J. Galbally, “Introduction to face presentation attack detection,” in *Handbook of Biometric Anti-Spoofing*, S. Marcel et al., Eds., pp. 187–206. Springer, 2019.
- [11] P. Tome, R. Vera-Rodriguez, J. Fierrez, and J. Ortega-Garcia, “Facial soft biometric features for forensic face recognition,” *Forensic Science International*, vol. 257, pp. 171–284, 2015.
- [12] K. Chang, K. W. Bowyer, and P. J. Flynn, “An evaluation of multimodal 2D + 3D face biometrics,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 27, pp. 619–624, 2005.
- [13] Y. Qian, W. Deng, and J. Hu, “Unsupervised face normalization with extreme pose and expression in the wild,” in *IEEE CVPR*, 2019.
- [14] P. Ekman, W. Friesen, and J. Hager, *The Facial Action Coding System*, Salt Lake City, UT, 2002.
- [15] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *ICLR*, 2015.
- [16] Q. Cao, L. Shen, et al., “VGGFace2: A dataset for recognising faces across pose and age,” in *IEEE FG*, 2018.
- [17] K. Zhang, Z. Zhang, et al., “Joint face detection and alignment using multitask cascaded convolutional networks,” *IEEE Signal Processing Letters*, vol. 23, no. 10, pp. 1499–1503, 2016.
- [18] J. Deng, J. Guo, et al., “ArcFace: Additive angular margin loss for deep face recognition,” in *IEEE CVPR*, 2019.
- [19] D. Han, J. Kim, and J. Kim, “Deep pyramidal residual networks,” in *IEEE CVPR*, 2017.
- [20] J. Deng, J. Guo, Ververas, et al., “RetinaFace: Single-shot multi-level face localisation in the wild,” in *IEEE CVPR*, 2020.
- [21] S. Du, Y. Tao, and A. M. Martinez, “Compound facial expressions of emotion,” *Proc. of the National Academy of Sciences*, vol. 111, pp. 1454–1462, 2014.
- [22] P. Lucey, J. F. Cohn, et al., “The extended Cohn-Kanade dataset (CK+): A complete dataset for action unit and emotion-specified expression,” in *IEEE CVPRW*, 2010.
- [23] Z. Liu, P. Luo, X. Wang, and X. Tang, “Deep learning face attributes in the wild,” in *ICCV*, 2015.
- [24] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao, “MS-Celeb-1M: A dataset and benchmark for large-scale face recognition,” in *IEEE ECCV*, 2016.
- [25] I. Kemelmacher-Shlizerman, S. M. Seitz, D. Miller, and E. Brossard, “The MegaFace benchmark: 1 million faces for recognition at scale,” in *IEEE CVPR*, 2016.
- [26] B. Maze, J. Adams, et al., “IARPA Janus Benchmark - C: Face Dataset and Protocol,” in *IAPR ICB*, 2018.
- [27] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, “Labeled faces in the wild: A database for studying face recognition in unconstrained environments,” Tech. Rep. 07-49, University of Massachusetts, Amherst, October 2007.
- [28] I. Serna, A. Morales, J. Fierrez, M. Cebrian, N. Obradovich, and I. Rahwan, “SensitiveLoss: Improving accuracy and fairness of face representations with discrimination-aware deep learning,” *arXiv/2004.11246*, 2020.
- [29] L. Kulke et al., “A comparison of the Affectiva iMotions facial expression analysis software with EMG for identifying facial expressions of emotion,” *Frontiers in Psychology*, 2020.
- [30] S. Gong, X. Liu, and A. K. Jain, “Jointly de-biasing face recognition and demographic attribute estimation,” in *IEEE ECCV*, 2020, pp. 330–347.
- [31] A. Morales, J. Fierrez, R. Vera-Rodriguez, and R. Tolosana, “SensitiveNets: Learning agnostic representations with application to face recognition,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 43, pp. 2158–2164, 2021.
- [32] P. Terhorst et al., “A comprehensive study on face recognition biases beyond demographics,” *arXiv/2103.01592*, 2021.
- [33] Y. Nirkin, Y. Keller, and T. Hassner, “FSGAN: Subject agnostic face swapping and reenactment,” in *ICCV*, 2019.
- [34] R. Tolosana, R. Vera-Rodriguez, J. Fierrez, et al., “DeepFakes and beyond: A survey of face manipulation and fake detection,” *Information Fusion*, vol. 64, pp. 131–148, 2020.
- [35] J.C. Neves et al., “GANprintR: Improved fakes and evaluation of the state of the art in face manipulation detection,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, no. 5, pp. 1038–1048, August 2020.
- [36] P. Perera, J. Fierrez, and V. Patel, “Quickest intruder detection for multiple user active authentication,” in *IEEE Intl. Conf. on Image Processing (ICIP)*, October 2020.