



Universidad Autónoma
de Madrid

Biblos-e Archivo
Repositorio Institucional UAM

Repositorio Institucional de la Universidad Autónoma de Madrid
<https://repositorio.uam.es>

Esta es la **versión de autor** del artículo publicado en:
This is an **author produced version** of a paper published in:

A. Peña, J. Fierrez, A. Morales and A. Lapedriza, "Learning Emotional-Blinded Face Representations," *2020 25th International Conference on Pattern Recognition (ICPR)*, 2021, pp. 3566-3573

DOI: <http://doi.org/10.1109/ICPR48806.2021.9412581>

Copyright: © 2020 Institute of Electrical and Electronics Engineers

El acceso a la versión del editor puede requerir la suscripción del recurso
Access to the published version may require subscription

Learning Emotional-Blinded Face Representations

Alejandro Peña, Julian Fierrez, Aythami Morales

School of Engineering

Universidad Autonoma de Madrid, Madrid, Spain

{alejandro.penna, julian.fierrez, aythami.morales}@uam.es

Agata Lapedriza

Universitat Oberta de Catalunya, Barcelona, Spain

Massachusetts Institute of Technology, Cambridge, USA

alapedriza@uoc.edu/agata@mit.edu

Abstract—We propose two face representations that are blind to facial expressions associated to emotional responses. This work is in part motivated by new international regulations for personal data protection, which enforce data controllers to protect any kind of sensitive information involved in automatic processes. The advances in Affective Computing have contributed to improve human-machine interfaces but, at the same time, the capacity to monitorize emotional responses triggers potential risks for humans, both in terms of fairness and privacy. We propose two different methods to learn these expression-blinded facial features. We show that it is possible to eliminate information related to emotion recognition tasks, while the performance of subject verification, gender recognition, and ethnicity classification are just slightly affected. We also present an application to train fairer classifiers in a case study of attractiveness classification with respect to a protected facial expression attribute. The results demonstrate that it is possible to reduce emotional information in the face representation while retaining competitive performance in other face-based artificial intelligence tasks.

I. INTRODUCTION

During the past 15 years there has been a lot of effort in creating technologies to extract emotional information from facial expressions [1, 2]. These facial analysis technologies can contribute to improve human-centric AI applications, like enhancing the user experience [3] or facilitating the human-computer interaction [4].

However, with the increase of image-capturing devices and available software for face image processing, face analysis technologies can also trigger potential risks for humans, both in terms of fairness and privacy. First, facial analysis software inherits human biases [5, 6], making them to perform poorly or unfairly on groups of population that are not well represented in the training data [7]. Second, humans might want to keep their emotions private or to make sure emotion recognition software is not used without their consent. Notice that privacy protection is deeply embedded in the normative framework that underlies various national and international regulations. For example, in April 2018 the European Parliament adopted a set of laws aimed to regularize the collection, storage and use of personal information [8]. In particular, these laws encourage to integrate privacy preserving methods in the technology when it is created.

As a possible solution for preserving the users privacy in the context of automatic face recognition, we propose to extract face features that are blind to facial expressions. As shown in Sec. V-A, generic face features learned for the task of subject recognition preserve information to perform tasks related to

facial expression classification. However, features extracted for the target task of subject recognition do not need to preserve this facial expression information. In this paper we show that we can effectively learn alternative face feature representations for the task of subject classification that are blind to facial expression. Notice that our work is in the direction of creating automatic emotion-suppression systems, i.e., algorithms to automatically remove emotional information from captured data, with the goal of preserving privacy. A similar idea was recently explored in [9], where the goal is to suppress physiological information from facial videos. Both facial expressions and physiological signals contain information related to emotional states.

In Sec. III-A we formally describe the problem of learning the emotional-blinded face representations. Then, we propose two different methods to learn these expression-blinded facial features, which are based on existing generic techniques for learning agnostic representations. The first one (*SensitiveNets*) consists of learning a discriminator for the target task and at the same time an adversarial regularizer to reduce facial expression information. The second one (*Learning not to Learn*) consists of using a regularized loss function during learning, which quantifies the amount of information on the sensitive task (facial expression recognition) by computing the mutual information between the feature space and a pre-trained facial expression classifier. The details of these two methods can be found in Sec. III-B.

To validate the proposed framework and methods we perform an extensive set of experiments (Sec. V). First, we show that face features learned for subject verification contain significant information to perform facial expression classification (sensitive information). Then, we show that both of the proposed methods can actually eliminate information related to facial expression. In particular, for the first method, we show how the facial expression recognition accuracy drops significantly when our proposed blinded face representations are applied, while the performance of subject verification, gender recognition, and ethnicity classification are just slightly affected. Finally, our last experiment shows how the proposed methods can be applied in another face analysis problem (Attractiveness Classification) to protect the emotional information.



Fig. 1. Visual examples of the facial expressions, corresponding to basic emotions, that are used in our experiments. Images from CFEE database [10].

II. RELATED WORKS

The study of new learned representations to improve the fairness of learning processes has attracted the attention of researchers [11–14]. In particular, [12, 13] proposed projection methods to preserve individual information while obfuscating membership to specific groups. The main drawback of these techniques was that discrimination was modelled as statistical imparity, which is not applicable when the classification task does not correlate with membership in a specific group.

Bias correction and sensitive information removal are related to each other but they are not necessarily the same thing. Bias is traditionally associated with unequal representation of classes in a dataset [15]. Dataset bias can produce unwanted results in the decision-making of algorithms, e.g., different face recognition accuracy depending of your ethnicity [16, 17]. Researchers have explored new learning processes capable to compensate this dataset bias [18, 19], but the correction of biased training processes does not necessarily serve to eliminate sensitive information from the trained representation. While the correction of biased models seeks to generate representations that perform similarly for different groups or classes, the removal of sensitive information seeks to eliminate this information from that representation. The proposal in [18] is based on a joint learning and unlearning algorithm inspired in domain and task adaptation methods [20]. The authors of [21] propose a new regularization loss based on mutual information between feature embeddings and bias, training the networks using adversarial [22] and gradient reversal [23] techniques. Finally, in [24] a privacy-preserving learning method is proposed to remove sensitive information in feature embeddings, without losing performance in the main task. These works reported encouraging results showing that it is possible to remove sensitive information (named as spurious variations in [18]) for age, gender, ancestral origin, and pose in face processing for different applications [25].

On the other hand, the normalization of face images directly in the raw image space according to specific face attributes such as pose [26] or gender [27, 28] is a challenging task. In [27] researchers proposed de-identification techniques that obfuscate gender attributes while preserving face verification

accuracy. Similarly, the method was based on Generative Adversarial Networks trained to generate androgynous images capable of fooling gender detection systems. The method in [26] proposed 3D models to normalize the face expressions. Although these methods showed promising results to generate realistic images, the main drawback of these techniques is that sensitive information is not eliminated but distorted. In [24], researchers demonstrated that sensitive information can be easily detected in those images when supervised learning processes are trained in the distorted domain.

A. How Emotions are Expressed in Face Images

Automatic emotion perception from facial expressions is an active area of research [29]. Some methods are based on the Facial Action Coding System [30], which encodes the facial expression using a set of specific localized face movements, called Action Units (AU). State-of-the-art systems for AU detection consist of deep learning models trained with large datasets [31]. These methods show impressive accuracies, even in uncontrolled environments [32]. However, while there are systems for AUs detection that are accurate enough to be used in practical applications, the prediction of emotions from these face movements is a more challenging problem. In that case, given a specific configuration of these face movements (that we call facial expression) the goal is to recognize the emotion category expressed by the face. There are several works on face analysis that attempt to recognize the 6 basic emotions proposed by Ekman and Friesen [33] [2] or emotional dimensions, such as *valence*, *arousal*, and *dominance* [29]. In general, all these methods are partially based on the assumption that each emotion is universally expressed with a specific face movement or, equivalently, with a specific combinations of AUs (see Fig. 1).

On the contrary, there are studies showing that there is no universal correspondence between AUs and emotions and, therefore, it is not always possible to recognize emotions just with the information provided by facial expressions [34]. Although this lack of agreement on whether it is possible or not, in certain circumstances, to recognize emotions just from facial expressions, the studies on psychology consistently show that facial movements and expressions communicate a lot of information, including information related to emotional states [34, 35]. Thus, learning face features that are blind to facial expressions, as proposed in this paper, can actually contribute to preserve emotion privacy.

Additionally, understanding how facial expressions are represented in feature embeddings of deep neural networks models is important to gain insights into the learning processes of these algorithms. Most face recognition algorithms are trained to be agnostic to this information (i.e. facial expressions may change and these changes should not affect the recognition tasks). However, the features used to recognize a face are also useful in general to recognize face gestures. Face expression databases traditionally include both AUs and emotion labels [36]. These databases are usually employed to model face gestures as well as affective interfaces.

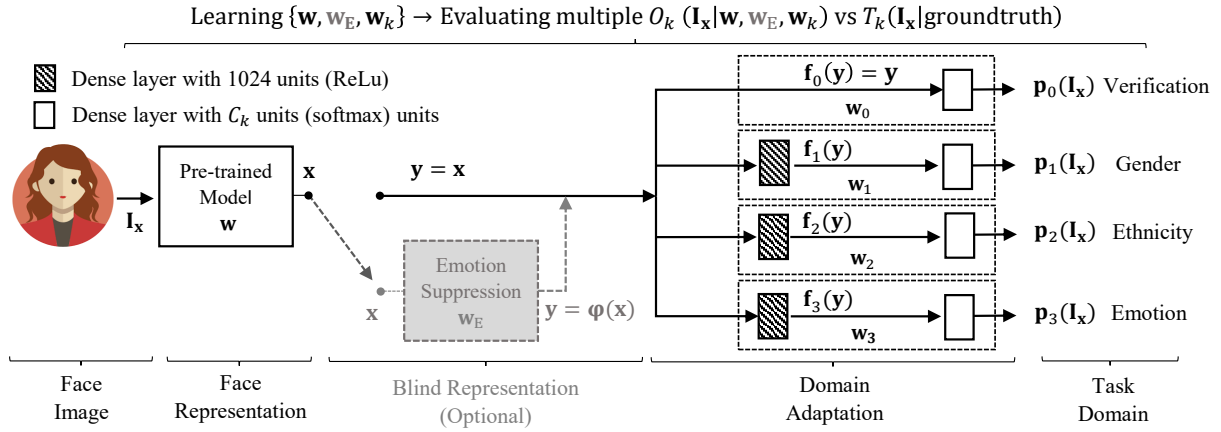


Fig. 2. General framework including domain adaptation from a pre-trained face representation to multiple tasks ($k = 0$ to 3) with and without the emotional-blinded representation $\varphi(\mathbf{x})$. C_k is the number of classes for task k (e.g. $C_1 = 2$ corresponds to classes *male* and *female*). \mathbf{f}_k is the projection to the adapted domain and \mathbf{p}_k is a vector with the probabilities of each class for task k .

III. LEARNING EMOTIONAL-BLINDED REPRESENTATIONS

A. Problem Formulation

We employ the privacy-preserving learning framework showed in Fig. 2 and detailed in [24]. The feature vector $\mathbf{x} \in \mathbb{R}^N$ is a face representation (a.k.a face embedding) obtained as the output of one of the last layers of a face model defined by parameters $\mathbf{w} \in \mathbb{R}^M$. In our framework, the parameters of the model \mathbf{w} were trained to reveal patterns associated to the identity of face images \mathbf{I}_x (i.e. face verification). This pre-trained model and the databases employed for training will be detailed in Sec. V-A.

In this framework, domain adaptation is used to transform the original representation trained for face verification $\mathbf{f}_0(\mathbf{y})$ into a new representation $\mathbf{f}_k(\mathbf{y})$ ($k \geq 1$ in Fig. 2) trained for different tasks ($k = 1$: for Gender classification, $k = 2$: Ethnicity classification, $k = 3$: Emotion classification). This adaptation is performed leaving fixed $\mathbf{w} = \mathbf{w}^*$ as obtained in the pre-trained model. The domain adaptation process for a task $k \geq 1$ results in a new learned model \mathbf{w}_k^* used to transform the original representation \mathbf{x} for the specific task k . Given a face image, the final output of the learned model (i.e. pre-trained plus domain adaptation) is a vector $\mathbf{p}_k(\mathbf{I}_x)$ containing C_k probabilities associated to each of the classes of task k .

In this work, we evaluate the face embeddings generated by the pre-trained model according to its performance in the original task (i.e. face verification) and 3 other different tasks: 1) Gender Classification; 2) Ethnicity Classification; and 3) Emotion Classification based on five of the six basic emotions proposed by Ekman plus the neutral expression (*Neutral, Happy, Sad, Disgusted, Angry, Surprised*).

The models $\{\mathbf{w}^*, \mathbf{w}_k^*\}$ are trained for a given task k represented by its target function T_k . The aim of the learning process is to minimize the error between the output O_k of the model and the target function T_k (e.g. $T_1 = 1$ for *male* and $T_1 = 0$ for *female*). The most popular approach for that is to train \mathbf{w} and \mathbf{w}_k by minimizing a loss function \mathcal{L}_1 over a

set of Pre-training samples \mathcal{P} for which we have groundtruth targets:

$$\min_{\mathbf{w}, \mathbf{w}_k} \sum_{\mathbf{I}_x \in \mathcal{P}} \mathcal{L}_1 [O_k(\mathbf{I}_x|\mathbf{w}, \mathbf{w}_k), T_k(\mathbf{I}_x|\text{groundtruth})] \quad (1)$$

The parameters $\{\mathbf{w}^*, \mathbf{w}_k^*\}$, trained using Eq. (1), generate a representation $\mathbf{f}_k(\mathbf{y})$ that maximizes the performance of the model for the task k .

In this framework, the goal of emotional-blinded learning starting from pre-trained networks is to solve, including or not the Emotion Suppression module, the following problem:

$$\min_{\mathbf{w}, \mathbf{w}_E, \mathbf{w}_k} \sum_{\mathbf{I}_x \in \mathcal{S}} \{ \mathcal{L}_1 [O_k(\mathbf{I}_x|\mathbf{w}, \mathbf{w}_E, \mathbf{w}_k), T_k(\mathbf{I}_x|\text{groundtruth})] + \mathcal{L}_2 [O_3(\mathbf{I}_x|\mathbf{w}, \mathbf{w}_E, \mathbf{w}_3), T_3(\mathbf{I}_x|\text{groundtruth})] \} \quad (2)$$

where \mathcal{L}_2 represents a loss function intended to minimize performance in the emotion recognition task T_3 while \mathcal{L}_1 tries to maximize performance in a different task T_k . In our experiments we use T_0 (Face Verification) as a task to maximize the performance. In the optimization problem (2) we may use a Suppression training dataset \mathcal{S} different to \mathcal{P} , and the optimization can take advantage of a previous solution $\{\mathbf{w}^*, \mathbf{w}_k^*\}$ to (1) in different ways. Let's denote the solution to (2) as $\{\mathbf{w}^{**}, \mathbf{w}_E^{**}, \mathbf{w}_k^{**}\}$.

In our experiments, we begin without Emotion Suppression ($\mathbf{y} = \mathbf{x}$ in Fig. 2) generating \mathbf{w}^* in a face recognition task by pre-training using the VGGFace2 database (3 million images from more than 9,000 people [37]). We then fix \mathbf{w}^* and train the Emotion classifier \mathbf{w}_3^* with the CFEE database (1,380 images from 230 people, with 6 images per subject, corresponding each of these 6 images to a different emotion [10]). Finally, we solve Eq. (2) considering $\{\mathbf{w}^*, \mathbf{w}_k^*\}$ as a starting point for obtaining the solution $\{\mathbf{w}^{**}, \mathbf{w}_E^{**}, \mathbf{w}_k^{**}\}$ taking various optimization shortcuts as detailed in the following.

B. Suppressing Emotions from Face Representations

1) *Method 1 - SensitiveNets*: The work [24] recently proposed a general method to generate privacy-preserving representations starting from pre-trained networks. Here we adapt that approach to remove emotional information for the primary task k from 0 to 2 in Fig. 2.

Applying SensitiveNets to the general methodology presented before leads to: 1) fixing $\mathbf{w}^{**} = \mathbf{w}^*$, 2) activating the Emotion Suppression block $\varphi_{\text{SN}}(\mathbf{x})$ (SN for SensitiveNets) in Fig. 2, and then 3) solving the following version of Eq. (2):

$$\begin{aligned} \min_{\mathbf{w}_E, \mathbf{w}_3} \sum_{\text{triplet} \in \mathcal{S}_P} \{ \mathcal{L}_1[O_k(\text{triplet}|\mathbf{w}_E, \mathbf{w}_3), T_k(\text{triplet}|\text{groundtruth})] + \\ + \Delta^A + \Delta^P + \Delta^N \} \\ \text{s.t. } \max_{\text{triplet} \in \mathcal{S}_E} \text{Performance}_{k=3}(\varphi_{\text{SN}}(\mathbf{x}_{\text{triplet}}|\mathbf{w}_E), \mathbf{w}_3) \end{aligned} \quad (3)$$

where $\text{triplet} = \{\mathbf{I}_A, \mathbf{I}_P, \mathbf{I}_N\}$, \mathbf{I}_A and \mathbf{I}_P are face images of the same person, \mathbf{I}_N is a face image of a different person, \mathcal{L}_1 is the triplet loss function proposed for face recognition in [38][39], and the three Δ terms are adversarial regularizers used to measure the amount of emotion information in the learned model represented by \mathbf{w}_E :

$$\Delta = \log\{1 + |0.9 - P_3(\text{Neutral} | \varphi_{\text{SN}}(\mathbf{x}|\mathbf{w}_E), \mathbf{w}_3)|\} \quad (4)$$

The probability P_3 of observing a *Neutral* expression in the face embedding after Emotion Suppression (φ_{SN}) is initially obtained with the pre-trained Emotion classifier \mathbf{w}_3^* , and SensitiveNets then iterates to solve Eq. (3) in order to obtain \mathbf{w}_E^{**} (the Emotional Suppression projection) and \mathbf{w}_3^{**} (an adapted Emotion classifier). In Eq. (4) $|\cdot|$ is the absolute value, and the Δ terms will tend to zero for larger P_3 . Therefore, by minimizing them in Eq. (3) we force the training to output *Neutral* expression in general, in this way eliminating the capacity to detect expressions other than *Neutral* from the face representation $\varphi_{\text{SN}}(\mathbf{x})$. In other words, we unlearn the facial features necessary to differentiate between different expressions.

On the other hand, Eq. (3) includes a constraint that will be enforced in subsequent iterations of SensitiveNets in a kind of min-max adversarial formulation [40]. Eq. (3) thus minimizes the emotion information in $\varphi_{\text{SN}}(\mathbf{x})$ with the Δ terms, trying to classify emotions based on $\varphi_{\text{SN}}(\mathbf{x})$ in the iterative learning with the optimization constraint (with decreasing success as the learning progresses), and maintaining the performance in the primary task with the triplet loss term \mathcal{L}_1 .

For solving Eq. (3) we apply the iterative adversarial learning approach proposed in [24] using the CFEE database [10] as \mathcal{S}_E to retrain the emotion detector (i.e., enforcing the constraint), and the DiveFace database [24] as \mathcal{S}_P to maintain the recognition accuracy.

The network \mathbf{w}_E consists of three dense layers with 1024 units each layer (linear activation). After solving Eq. (3) the network \mathbf{w}_E^{**} generates the emotional blinded representation

$\varphi_{\text{SN}}(\mathbf{x})$, which removes sensitive information (emotions in the present paper) while maintaining recognition performances.

2) *Method 2 - Learning not to Learn*: The second approach studied here to remove emotional features is based on [21]. Similar to SensitiveNets [24], this method uses a regularization algorithm to train deep neural networks, in order to prevent them from learning a known factor present in the training set irrelevant or undesired for a given primary task. Here we propose to unlearn emotional features for the primary task k from 0 to 2 in Fig. 2.

In this case the Emotion Suppression switch is off, therefore there is no \mathbf{w}_E , and we start from pre-trained $\{\mathbf{w}^*, \mathbf{w}_k^*, \mathbf{w}_3^*\}$.

The training algorithm uses a regularization loss that includes the mutual information between emotions and feature embeddings \mathbf{x} . These embeddings are then fed into both the main classification task network (corresponding to k from 0 to 2), and the emotion classification network \mathbf{p}_3 . The function to optimize for emotion removal is then:

$$\begin{aligned} \min_{\mathbf{w}, \mathbf{w}_k} \sum_{\mathbf{I}_x \in \mathcal{S}} \{ \mathcal{L}_c[O_k(\mathbf{I}_x|\mathbf{w}, \mathbf{w}_k), T_k(\mathbf{I}_x|\text{groundtruth})] + \\ + \lambda \mathcal{I}[\mathbf{p}_3(\mathbf{I}_x); \mathbf{x}] \} \end{aligned} \quad (5)$$

where \mathcal{L}_c denotes the cross-entropy loss, \mathcal{I} represents the mutual information and λ is an hyper-parameter.

To compute the mutual information in Eq. (5), we used the emotion classification network to approximate the a posteriori distribution of the emotional classifier $\mathbf{p}_3(\mathbf{I}_x)$. The training algorithm can be implemented in practice following an adversarial strategy [22], combined with the use of the gradient reversal technique [23].

IV. DATA AND EXPERIMENTAL SET UP

To obtain the face representation \mathbf{x} we use a learning architecture with state-of-the-art performance in face recognition tasks: ResNet50, proposed in [41]. ResNet50 has around 41M parameters split in 34 residual layers. The pre-trained model used in this work was trained from scratch with VGGface2 dataset [37]. This ResNet50 model achieved 98.0% accuracy in face verification with the IJB-A dataset [42].

Using the base representation \mathbf{x} generated by the pre-trained network ResNet50 we trained different classifiers as depicted in Fig. 2 according to the following labeled databases:

- DiveFace [24]: The DiveFace database contains annotations equitably distributed among 6 demographic classes, related to gender and 3 ethnic groups (*East Asian* | *Sub-Saharan* and *South Indian* | *Caucasian*), with 24K different identities and a minimum of 3 images per identity. This database was used to train the emotional-blinded representation. Additionally, 12K subjects of this database were used to train and test the gender and ethnicity classification.
- CFEE [10]: The Compound Facial Expressions of Emotion database includes facial images of 230 different users. For every user, we selected an image belonging to

each of the 22 categories present in the dataset: 6 basic emotions, 15 compound emotions (i.e. a combination of two basic emotion), and neutral expression. All images represent a fully recognizable expression, being captured in a controlled environment of illumination and pose. We used the 6 basic emotion of this database to train the emotional-blinded representation.

- LFW [43]: Labeled Faces in the Wild is a database for research on unconstrained face recognition. It contains more than 13K images of faces collected from the web. We employ the aligned images [44] from the test set provided with view 1 and its associated evaluation protocol.
- CelebA [45]: The CelebA dataset has a total of 202K celebrity images from more than 10K identities. Each image is annotated with 40 binary attributes, including appearance features, gender, age, attractiveness and emotional state, and 5 landmark positions. The dataset is partitioned into 2 splits, with 8K identities retained as the training set, and the remaining 2K as the test set.

In order to measure how much emotional information is available in the face representation, we trained different emotion classifiers using either original embeddings \mathbf{x} or emotional-blinded representations $\varphi(\mathbf{x})$. We measured the amount of emotional information as the performance achieved by these classification algorithms. We assume that emotional information is removed by our blinding transformation $\varphi(\cdot)$ when a significantly drop of performance in emotion classification occurs in comparison to the original emotion classification accuracy before applying that transformation.

The face recognition accuracy is obtained according to the evaluation protocol of the popular benchmark of LFW [43]. For the rest of tasks, we used 80% of the samples for training and 20% for testing. Implementation details: 150 epochs, Adam optimizer (learning rate = 0.001, $\beta_1 = 0.9$, and $\beta_2 = 0.999$), and batch size of 128 samples.

V. EXPERIMENTS

A. Are Facial Expressions Encoded in Generic Face Representations?

To better understand how the emotional features are embedded in the deep face representations, we study how identity and emotional information are represented in \mathbf{x} and $\mathbf{f}_3(\mathbf{x})$.

Fig. 3 shows the two-dimensional t-SNE projection of the original face representation \mathbf{x} and the learned representation $\mathbf{f}_3(\mathbf{x})$ for emotion recognition using the CFEE database [10] (detailed in Sec. IV). This database is interesting for this study because of its controlled acquisition environment (covariates such as pose or illumination are not present) and the multiple face gestures available for 230 subjects. We ran t-SNE over \mathbf{x} and $\mathbf{f}_3(\mathbf{x})$ without using the emotion labels available, and then show in Fig. 3 the resulting t-SNE projections with emotion labels *a posteriori* for visualization purposes.

As we can see in Fig. 3 (top), the projection in the original representation ignores the emotional features. The representation learned for face verification deprecates emotional features

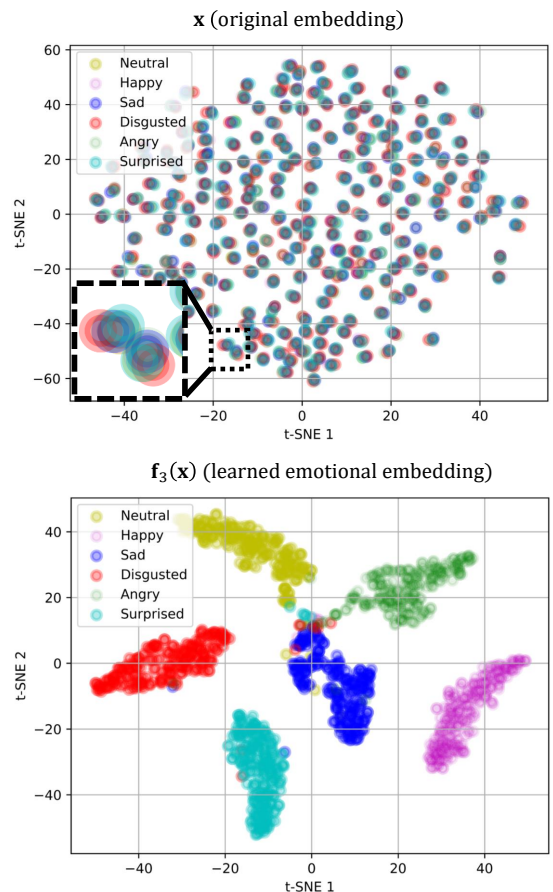


Fig. 3. t-SNE plot of the original embedding \mathbf{x} (top) and emotion feature transformation $\mathbf{f}_3(\mathbf{x})$ (bottom) of the face images from CFEE database.

in order to maximize accuracy in face recognition. Face expressions can be seen as distortions that should be removed from the decision-making of the representation. However, if we freeze the weights of the ResNet model that produced the representation \mathbf{x} and we train the representation $\mathbf{f}_3(\mathbf{x})$ for Emotion Classification, we can observe in Fig. 3 (bottom) how emotional features were available in \mathbf{x} and a simple training procedure with hundreds of samples allows to extract that information and correctly classify the emotions for more than 90% of the face images. Note that as mentioned before, ResNet was trained originally for identity recognition and these emotional features were not intentionally included in the learning process. These results illustrate that emotional information is embedded in \mathbf{x} even though that representation was trained for a different purpose (i.e. face verification).

To gain insight into how the emotional features are embedded in the original representation \mathbf{x} , we have evaluated the performance of an emotion classifier when different amount of features from \mathbf{x} are available to train $\mathbf{f}_3(\mathbf{x})$. To do this, in each iteration we randomly suppress a percentage of features of the representation \mathbf{x} and we re-train the emotion representation $\mathbf{f}_3(\mathbf{x})$, always freezing the ResNet model. Fig. 4 shows the performance decay for Emotion Classification related to the




Fig. 4. Performance of the emotion classifier $p_3(\mathbf{I}_x)$ vs number of features from \mathbf{x} used to train \mathbf{w}_3 with fixed \mathbf{w} (see Fig. 2).

number of features suppressed from the original representation \mathbf{x} . It is remarkable how well the emotion representation is capable of classifying with 70% accuracy even if the number of features available is only 10% of the original size. The model is able to keep almost the same performance until 90% of features are suppressed. This demonstrates that emotional features are latent in almost all features of the original representation \mathbf{x} .

B. Emotional-Blinded Face Representations

The goal is to keep the recognition capability in other face classification tasks while removing the emotion information embedded in the face representation \mathbf{x} using the methods described in Sec. III-B for generating $\phi(\mathbf{x})$ (see Fig. 2). To analyze the effectiveness of those blinding methods, we conducted experiments on 3 datasets (see Sec. IV): DiveFace, CFEE, and LFW.

a) *Objective 1 - Maintaining face representation information:* the goal is to maintain the performance of the emotional-blinded face representation for other tasks different to emotion classification. We calculated the performance of 3 different face-based machine learning tasks using either original embeddings \mathbf{x} or their projections $\phi(\mathbf{x})$. The tasks are evaluated according to the classification accuracy obtained in the test set. Table I shows the classification accuracy of representations generated by the pre-trained model before and after the projections $\phi_{SN}(\mathbf{x})$ obtained by the *Method 1* and $\phi_{LnL}(\mathbf{x})$ obtained using the *Method 2* (see Sec. III-B). The results of the projection $\phi_{SN}(\mathbf{x})$ show a very small drop of performance when the projection is applied in the first domains (ID, Gender, and Ethnicity), which demonstrates the success of our method in preserving most of the discriminative information in the face representation. The drop of performance in the method based on the $\phi_{LnL}(\mathbf{x})$ projection is higher for the primary tasks (ID, Gender, Ethnicity) and the emotion classification. This decay may be caused because of the disentanglement of primary and secondary tasks managed by the mutual information regularizer in Eq. (5). The method proposed in [21] was originally evaluated for problems with limited number of classes and face recognition requires feature spaces capable of allocating large number of classes (one per identity).

TABLE I

ACCURACY OF DIFFERENT CLASSIFIERS TRAINED WITH \mathbf{x} (BEFORE) OR $\phi(\mathbf{x})$ (AFTER). DIFF IS THE ACCURACY DROP RELATIVE TO RANDOM CHOICE (DIFF=100% REPRESENTS A RANDOM CHOICE CLASSIFIER):
DIFF = (BEFORE - AFTER)/(BEFORE - RANDOM CHOICE)

Information Domain	\mathbf{x}	$\phi_{SN}(\mathbf{x})$	Diff. SN	$\phi_{LnL}(\mathbf{x})$	Diff. LnL
ID	96.8	96.3	↓ 1%	59.4	↓ 75.0%
Gender	99.2	98.9	↓ 1%	72.7	↓ 53.9%
Ethnicity	98.8	98.6	↓ 1%	67.4	↓ 47.9%
Emotion (NN)	88.1	59.6	↓ 40%	41.6	↓ 65.0%
Emotion (SVM)	88.1	16.7	↓ 100%	25.0	↓ 88.2%
Emotion (RF)	77.4	58.3	↓ 31%	44.7	↓ 53.8%

b) *Objective 2 - Removing emotional information:* to analyze the amount of emotional information available in the face representations we train different emotion classification algorithms (NN = Neural Networks, SVM = Support Vector Machines, and RF = Random Forests) either on original embeddings \mathbf{x} or on their projections $\phi(\mathbf{x})$. Table I shows the accuracies obtained by each algorithm before and after the projections. Results show a significant drop of performance in classification when both blinding representations are applied, which demonstrates the success in reducing the emotion information from the embeddings. However, the emotional information is deeply embedded in the representations, and to keep the performance of other tasks (first 3 rows of Table I) not all the emotion information was removed.

There are differences between the performances obtained by the two blinding methods. While $\phi_{SN}(\mathbf{x})$ maintains higher performance in the primary tasks (ID, Gender, and Ethnicity), the emotion suppression is higher in $\phi_{LnL}(\mathbf{x})$. This higher suppression obtained by $\phi_{LnL}(\mathbf{x})$ may be due to the weaker representations generated by this method which lead to worse performance in the primary tasks. However, the accuracy obtained for emotion classification using both methods (lower than 60% in all cases) may be low enough to prevent its unwanted exploitation. Emotion-related privacy is not fully granted, but clearly improved.

Fig. 5 shows the two-dimensional t-SNE projection similar to Fig. 3 (bottom) for the emotional blinded representation $\mathbf{f}_3(\phi_{SN}(\mathbf{x}))$. The results show how the domain adaptation training of \mathbf{w}_3 (see Fig. 2) was not able to find a representation capable of discriminating emotions in the learned representation $\phi_{SN}(\mathbf{x})$.

C. Blind Representations: Towards Equality of Opportunity

Inspired in the experiments performed in [7] for analyzing biases and achieving a specific fairness criterion, here we study how blind representations can improve the *Equality of Opportunity* [46]. For this purpose we introduce task $k = 4$: binary Attractiveness classification (*Attractive* | *Not Attractive*) based on an input face image \mathbf{I}_x .

In this experiment, the outcome of an Attractiveness classifier with input \mathbf{x} and parameters \mathbf{w}_4 given its positive class should be independent to the feature s we want to protect in terms of fairness. In our experiments, the protected attribute




Fig. 5. t-SNE plot of the emotional embedding \mathbf{f}_3 trained with the proposed emotional-blinded representation $\varphi_{\text{SN}}(\mathbf{x})$ as input (see Fig. 2) over the CFEE database. See previous Fig. 3 (bottom) for comparison.

TABLE II
RESULTS ON ATTRACTIVENESS CLASSIFICATION (ACC = ACCURACY).
EQUAL OPPORTUNITIES IS CALCULATED AS:
 $100 - (\text{TPR SMILING} - \text{TPR NOT SMILING})$.
TPR = TRUE POSITIVE RATE IN ATTRACTIVENESS CLASSIFICATION

Method (training)	Acc.	TPR Smil.	TPR Not Smil.	Eq. Opp.
\mathbf{x} (unbiased)	77.26%	84.55%	82.47%	97.93%
\mathbf{x} (biased)	76.23%	84.17%	66.70%	82.53%
$\varphi_{\text{SN}}(\mathbf{x})$ (biased)	74.50%	81.87%	73.58%	91.71%
$\varphi_{\text{LnL}}(\mathbf{x})$ (biased)	76.62%	86.97%	73.70%	86.73%

is a specific face gesture: smile. Therefore, in our case: $s \in \{\textit{Smiling}, \textit{Not Smiling}\}$. Using the framework presented in Sec. III-A summarized in Fig. 2, the Equality of Opportunity results in: $\mathbf{p}_4(\mathbf{I}_x | \mathbf{w}^*, \mathbf{w}_4^*, T = 1, s) = \mathbf{p}_4(\mathbf{I}_x | \mathbf{w}^*, \mathbf{w}_4^*, T = 1)$. This criterion implies equal True Positive Rates across the different face gestures defined by s and the Attractiveness classifier defined by the parameters $\mathbf{w}^*, \mathbf{w}_4^*$.

We used 40K images from CelebA dataset [45], previously introduced in Sec. IV, to train the Attractiveness classifier. Since some studies suggest that face expressions, such as smile, can affect the perception of attractiveness, we specifically train a biased classifier. In particular, we employed the smiling annotation available in CelebA as a face gesture commonly associated to a positive emotion that can therefore introduce undesired bias. We generated an emotionally biased training set where the proportion of attractive people smiling and not smiling was 70% and 30% respectively. We introduced the opposite bias for the unattractive group with 30% and 70% of smiling and not smiling respectively. In order to avoid the appearance of other biases, we balanced the dataset in terms of attractiveness and gender, compensating the gender bias of the dataset (i.e. the proportion of attractive females is 67%, while for males is 27%). We also generated an unbiased dataset with 50% smiling and not smiling samples (randomly chosen and balanced with respect to gender).

The results in Table II show higher True Positive Rates (TPR) for the privileged class (*Smiling* in our experiment) in comparison with the non-privileged class (*Not Smiling*). The face gesture *Smiling* was irrelevant to classify the attractiveness (i.e. there was no correlation between the attributes *Smiling* and *Attractiveness*). However, a classifier trained on face embeddings \mathbf{x} generated by pre-trained models like ResNet50, tends to reproduce the bias introduced in the training datasets. Table II shows how the blind representations $\varphi_{\text{SN}}(\mathbf{x})$ and $\varphi_{\text{LnL}}(\mathbf{x})$ presented in Sec. III-B significantly reduce the gap between both classes by improving equality in 9% and 4% respectively. The blind representations avoid the network to exploit the latent variable related with the face gesture and reduce the impact of the biased training dataset.

Implementation details: the classifiers were composed by one fully connected layer (1024 units and ReLu activation) and one output unit (sigmoid activation), which we feed with face embeddings generated with the methods mentioned above. We repeated the experiment five times, using different training sets with 36K images from CelebA, and evaluating the resulting classifiers on validation sets with 4K images, selected from the CelebA's evaluation split.

VI. CONCLUSIONS

The growth of emotion recognition technologies has allowed great advances in fields related to human-machine interaction. At the same time, having automatic systems capable to read emotions without explicit consent triggers potential risks for humans, both in terms of fairness and privacy. In this work we have proposed two face representations that are blind to facial expressions associated to emotional responses.

In addition to a general formulation of the problem, we have adapted two existing methods for this purpose of generating emotional-blinded face representations: SensitiveNets [24] and Learning not to Learn [21]. The results show that it is possible to reduce dramatically the performance of emotion classifiers (more than 40%) while the performance in other face analysis tasks (verification, gender, and ethnicity recognition) is only slightly reduced (less than 2%).

Finally, we included an experiment on facial attractiveness classification to show how to treat facial expression as protected information in face classification problems. The results show how blinded representations can improve a specific fairness criterion based on the principles and methods studied in the present paper.

VII. ACKNOWLEDGMENTS

This work has been supported by projects: PRIMA (H2020-MSCA-ITN-2019-860315), TRESPASS-ETN (H2020-MSCA-ITN-2019-860813), IDEA-FAST (IMI2-2018-15-853981), BIBECA (RTI2018-101248-B-I00 MINECO/FEDER), REAVIPERO (RED2018-102511-T), RTI2018-095232-B-C22 MINECO, and Accenture. A. Pea is supported by a research fellowship (PEJ2018-004094A) from the Spanish MINECO.

REFERENCES

- [1] M. Pantic and L. J. Rothkrantz, "Expert system for automatic analysis of facial expressions," *Image and Vision Computing*, vol. 18, no. 11, pp. 881–905, 2000.
- [2] S. Li and W. Deng, "Deep facial expression recognition: A survey," *IEEE Transactions on Affective Computing*, vol. 2010, 2020.
- [3] F. Eyben, M. Willmer *et al.*, "Emotion on the roadnecesity, acceptance, and feasibility of affective computing in the car," *Advances in Human-Computer Interaction*, 10 2010.
- [4] S. S. Guillén, L. L. Iacono, and C. Meder, "Affective robots: Evaluation of automatic emotion recognition approaches on a humanoid robot towards emotionally intelligent machines," *Energy*, vol. 3643, 2018.
- [5] R. E. Jack, O. G. Garrod, H. Yu, R. Caldara, and P. G. Schyns, "Facial expressions of emotion are not culturally universal," *Proceedings of the National Academy of Sciences*, vol. 109, no. 19, pp. 7241–7244, 2012.
- [6] G. Bijlstra, R. W. Holland, R. Dotsch, K. Hugenberg, and D. H. Wigboldus, "Stereotype associations and emotion recognition," *Personality and Social Psychology Bulletin*, vol. 40, no. 5, pp. 567–577, 2014.
- [7] N. Quadrianto, V. Sharmanska, and O. Thomas, "Discovering fair representations in the data domain," in *IEEE Conference on Computer Vision and Pattern Recognition*, June 2019.
- [8] EU 2016/679 (General Data Protection Regulation). European Commission. [Online]. Available: <https://gdpr-info.eu/>
- [9] W. Chen and R. W. Picard, "Eliminating physiological information from facial videos," in *IEEE International Conference on Automatic Face & Gesture Recognition*, 2017, pp. 48–55.
- [10] S. Du, Y. Tao, and A. M. Martinez, "Compound facial expressions of emotion," *Proceedings of the National Academy of Sciences*, vol. 111, no. 15, pp. E1454–E1462, 2014.
- [11] S. Jia, T. Lansdall-Welfare, and N. Cristianini, "Right for the right reason: Training agnostic networks," in *Advances in Intelligent Data Analysis XVII*, W. Duivesteyn, A. Siebes, and A. Ukkonen, Eds., 2018, pp. 164–174.
- [12] E. Raff and J. Sylvester, "Gradient reversal against discrimination: A fair neural network learning approach," in *IEEE International Conference on Data Science and Advanced Analytics*, 2018, pp. 189–198.
- [13] R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork, "Learning fair representations," in *International Conference on Machine Learning*, 2013, p. 325333.
- [14] I. Serna, A. Morales, J. Fierrez, M. Cebrian, N. Obradovich, and I. Rahwan, "Algorithmic discrimination: Formulation and exploration in deep learning-based face biometrics," in *Proc. of AAAI Workshop on SafeAI*, Feb. 2020.
- [15] A. Torralba and A. A. Efros, "Unbiased look at dataset bias," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2011, pp. 1521–1528.
- [16] B. F. Klare *et al.*, "Face recognition performance: Role of demographic information," *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 6, pp. 1789–1801, 2012.
- [17] P. Drozdowski, C. Rathgeb, A. Dantcheva, N. Damer, and C. Busch, "Demographic Bias in Biometrics: A Survey on an Emerging Challenge," *IEEE Transactions on Technology and Society*, vol. 1, no. 2, pp. 89–103, 2020.
- [18] M. Alvi, A. Zisserman, and C. Nellaaker, "Turning a blind eye: Explicit removal of biases and variation from deep neural network embeddings," in *European Conference on Computer Vision Workshops*, 2018.
- [19] S. Nagpal, M. Singh, R. Singh, M. Vatsa, and N. Ratha, "Deep Learning for Face Recognition: Pride or Prejudiced?" *arXiv:1904.01219*, 2019.
- [20] E. Tzeng, J. Hoffman, T. Darrell, and K. Saenko, "Simultaneous deep transfer across domains and tasks," in *The IEEE International Conference on Computer Vision*, December 2015.
- [21] B. Kim, H. Kim, K. Kim, S. Kim, and J. Kim, "Learning not to learn: Training deep neural networks with biased data," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [22] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems 27*, 2014, pp. 2672–2680.
- [23] Y. Ganin, E. Ustinova *et al.*, "Domain-adversarial training of neural networks," *Journal of Machine Learning Research*, vol. 17, no. 1, p. 20962030, 2016.
- [24] A. Morales, J. Fierrez, R. Vera-Rodriguez, and R. Tolosana, "SensitiveNets: Learning agnostic representations with application to face images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [25] A. Pea, I. Serna, A. Morales, and J. Fierrez, "Bias in multimodal AI: Testbed for fair automatic recruitment," in *IEEE CVPR Workshop on Fair, Data Efficient and Trusted Computer Vision*, 2020.
- [26] Y. Qian, W. Deng, and J. Hu, "Unsupervised face normalization with extreme pose and expression in the wild," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [27] V. Mirjalili, S. Raschka, A. Namboodiri, and A. Ross, "Semi-adversarial networks: Convolutional autoencoders for imparting privacy to face images," in *International Conference on Biometrics*, 2018, pp. 82–89.
- [28] E. Gonzalez-Sosa, J. Fierrez, R. Vera-Rodriguez, and F. Alonso-Fernandez, "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.

- [29] J. Kossaifi, A. Toisoul, A. Bulat, Y. Panagakis, T. M. Hospedales, and M. Pantic, "Factorized higher-order cnns with an application to spatio-temporal emotion estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 6060–6069.
- [30] P. Ekman and W. V. Friesen, *Facial action coding system: a technique for the measurement of facial movement*. Palo Alto: Consulting Psychologists Press, 1978.
- [31] C. F. Benitez-Quiroz, R. Srinivasan, and A. M. Martinez, "Emotionet: An accurate, real-time algorithm for the automatic annotation of a million facial expressions in the wild," in *IEEE International Conference on Computer Vision & Pattern Recognition*, 2016.
- [32] C. F. Benitez-Quiroz, Y. Wang, and A. M. Martinez, "Recognition of action units in the wild with deep nets and a new global-local loss," in *International Conference on Computer Vision*, 2017, pp. 3990–3999.
- [33] P. Ekman and W. V. Friesen, "Constants across cultures in the face and emotion." *Journal of Personality and Social Psychology*, vol. 17, no. 2, p. 124, 1971.
- [34] L. F. Barrett, R. Adolphs, S. Marsella, A. M. Martinez, and S. D. Pollak, "Emotional expressions reconsidered: challenges to inferring emotion from human facial movements," *Psychological Science in the Public Interest*, vol. 20, no. 1, pp. 1–68, 2019.
- [35] J. A. DeVito, S. O'Rourke, and L. O'Neill, *Human Communication*. Longman, 2000.
- [36] C. Fabian Benitez-Quiroz, R. Srinivasan, and A. M. Martinez, "Emotionet: An accurate, real-time algorithm for the automatic annotation of a million facial expressions in the wild," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [37] 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.
- [38] O. M. Parkhi, A. Vedaldi, and A. Zisserman, "Deep face recognition," in *British Machine Vision Conference*, 2015.
- [39] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 815–823.
- [40] J. Wang, T. Zhang *et al.*, "Towards a unified min-max framework for adversarial exploration and robustness," *arXiv:1906.03563*, 2019.
- [41] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [42] B. F. Klare *et al.*, "Pushing the frontiers of unconstrained face detection and recognition: IARPA Janus Benchmark A," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 1931–1939.
- [43] 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," University of Massachusetts, Amherst, Tech. Rep. 07-49, October 2007.
- [44] G. B. Huang, M. Mattar, H. Lee, and E. Learned-Miller, "Learning to align from scratch," in *Advances in Neural Information Processing Systems*, 2012.
- [45] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *International Conference on Computer Vision*, 2015.
- [46] M. Hardt, E. Price, E. Price, and N. Srebro, "Equality of opportunity in supervised learning," in *Advances in Neural Information Processing Systems 29*, 2016.