

Generative Adversarial Stacked Autoencoders for Facial Pose Normalization and Emotion Recognition

Ariel Ruiz-Garcia ✉

School of Computing, Electronics and Mathematics
Coventry University
Coventry, UK
ariel.9garcia@gmail.com

Vasile Palade

Research Centre for Data Science
Coventry University
Coventry, UK
vasile.palade@coventry.ac.uk

Mark Elshaw

School of Computing, Electronics and Mathematics
Coventry University
Coventry, UK
mark.elshaw@coventry.ac.uk

Mariette Awad

Humans and Machines Lab
American University of Beirut
Beirut, Lebanon
ma162@aub.edu.lb

Abstract—In this work, we propose a novel Generative Adversarial Stacked Autoencoder that learns to map facial expressions with up to ± 60 degrees to an illumination invariant facial representation of 0 degrees. We accomplish this by using a novel convolutional layer that exploits both local and global spatial information, and a convolutional layer with a reduced number of parameters that exploits facial symmetry. Furthermore, we introduce a generative adversarial gradual greedy layer-wise learning algorithm designed to train Adversarial Autoencoders in an efficient and incremental manner. We demonstrate the efficiency of our method and report state-of-the-art performance on several facial emotion recognition corpora, including one collected in the wild.

Index Terms—Emotion Recognition, Facial Pose Normalization, Generative Adversarial Networks, Illumination Invariance, Generative Adversarial Stacked Autoencoders.

I. INTRODUCTION

Facial expression recognition continues to be of great interest in the machine learning community due to the many challenges it presents. Many works in the literature have proposed a variety of models that produce state-of-the-art accuracy on various facial expression corpora collected in controlled environments. However, most of these models are unable to deal with non-frontal facial expression images or with drastic changes in the environment, such as different lighting conditions. This is evident on models trained on data with nonuniform conditions collected in the wild [1], for which state-of-the-art is significantly lower than on datasets taken in controlled environments. This is partly due to the fact that in non-frontal faces —i.e. faces with pose greater than 0 degrees—much of the information essential for emotion recognition is nonexistent. Moreover, the more variations in facial pose the larger the data distribution, and the more difficult for a neural network to provide good generalization due to the high dimensional search space. In addition, for real-time emotion recognition in unconstrained environments, it is difficult to obtain images without facial pose. In this work we introduce

a model that explicitly addresses facial pose and illumination invariance. The proposed model can deal with faces with a facial pose of up to ± 60 degrees and several degrees of illumination. Our main contributions can be summarized as:

- a novel deep Generative Adversarial Stacked Convolutional Autoencoder model that learns to map faces with facial pose of up to ± 60 degrees to 0 degrees representations.
- a hybrid deep learning layer employing convolutional filters to retain spatial information and learn salient features, and fully connected units shared across the depth dimension to facilitate the reduction of facial pose.
- a convolutional layer with reduced number of parameter that exploits facial symmetry and learns from only one half of the face.
- a gradual greedy layer-wise algorithm for Generative Adversarial Autoencoders.
- an illumination and pose invariant emotion recognition classifier that produces state-of-the-art classification performance in images taken in both controlled and unconstrained environments.

We also show the difference of training for pose invariance only and for pose and illumination invariance. Our model is tested on several empirical facial expression datasets as well as on one collected in the wild.

II. RELATED WORKS

Deep neural networks are known to be difficult to train. This was particularly true before the introduction of several regularization techniques [2] in recent years, and before Batch Normalization [3]. Before such techniques existed, pre-training models was often a preferred choice for deep models over random initialization. Autoencoders [4] are suitable for such tasks given that they are trained in an unsupervised fashion, overcoming constraints imposed by lack of labelled data.

Autoencoders learn an encoder function f that maps an input image x to a hidden representation $h = f(x)$, and learn a function g that maps h to a reconstruction $y = g(f(x))$ where y is an approximation of x . However, in recent works, [5] it has been established that the target reconstruction does not need to be the same as the input x to the autoencoder. This is supported by the theory that to be useful, an autoencoder should only learn an approximation of the target reconstruction and not an identity function that replicates it [6].

In principle this establishes that the input and target vectors in an autoencoder do not need to be the same, and therefore we can learn a function that maps an input from a given distribution to a target that lies in a different distribution. Autoencoder can also be trained in a greedy layer-wise (GLW) [7] fashion which has proven to be more effective than joint training [8].

Generative Adversarial Autoencoders [9] are some of the latest autoencoder models that follow the trending popularity of Generative Adversarial Networks (GANs) [10]. GANs are composed of two networks: a generative model G and a discriminator model D . Both models are trained by playing a min-max adversarial game where the discriminator model tries to determine if a given sample is from the generator or the dataset. In contrast, the generator maps samples z from a prior distribution $p(z)$ and maps it to the data space. Generative Adversarial Autoencoders follow a similar approach where the generator is an autoencoder that maps an input x to a latent representation z that lies in an aggregate posterior distribution $q(z)$ and back to a reconstruction y which is an approximation of x . The discriminator network in this framework attempts to determine if a sample has been drawn from a prior distribution $p(z)$ or from the latent distribution $q(z)$.

Although GANs are mainly used for data synthesis, some works have explored their use in classification [11]. For emotion recognition, works employing GANs mainly focus on using GANs for data augmentations, whether in emotion recognition from speech [12] or from facial expressions [13], [14]. Most other works focus on generating data, but do not explore emotion recognition. Such works include multi-purpose face recognition [15], [16], or facial expression image completion [17].

Some works attempt to deal with some of the common challenges in emotion recognition, such as, illumination invariance. Contemporary attempts to address illumination invariance in the domain of facial expression recognition include the use of noise injection [4], blurring images with Gaussian filters [18], a combination of histograms, principal component analysis (PCA) and discrete cosine transforms [19], or complex models [18], [20] and very deep CNN architectures [21].

To the best of the knowledge of the authors, although many works target pose invariant face recognition, no existing work focuses on pose invariant facial expression recognition.

In this work we normalize facial pose φ , where $|\varphi| > 0$ to frontal images 0 degrees pose. We also normalize images with relative luminance Y to a target luminance μ . Considering that autoencoders allow us to learn a mapping from an input

image to a target image that does not necessarily lie in the same distribution, they are suitable for this task. This in effect means that we are interested in imposing a distribution on the input data to produce reconstructions that resemble the desired target. Adversarial autoencoders can facilitate this task as they uniformly impose a data distribution on the code vector, i.e. the hidden representation produced by the encoder element, to generate realistic reconstructions. Moreover, adversarial autoencoders are designed to produce very realistic reconstructions.

III. GENERATIVE ADVERSARIAL STACKED CONVOLUTIONAL AUTOENCODERS




Fig. 1: Visualization of the first shallow autoencoder in the GASCA model.

We introduce the Generative Adversarial Stacked Convolutional Autoencoders (GASCA) framework. Just like in conventional adversarial autoencoders, in a GASCA model, the discriminator attempts to tell whether a sample comes from the training dataset or if it is a reconstruction produced by the autoencoder.

Let x_φ be a sample from the data distribution $p_d(x_\varphi)$ and x_μ the sample from the data distribution $p_d(x_\mu)$ used as the desired target reconstruction. The autoencoder G model learns to map x_φ to a latent space z , note that this is not an aggregate posterior as in conventional adversarial autoencoders, and back to a reconstruction y that resembles x_μ and lies in the distribution $q(y)$. The discriminator D attempts to differentiate between y and x_μ .

In the conventional adversarial autoencoder framework [9] a distribution $p(z)$ —often a Gaussian distribution—is imposed on $q(z)$ by estimating the divergence between q and p . This imposition can be used to produce reconstructions with specific features. However, in this work, the objective is to produce reconstructions that are as close as possible to the desired target image x_μ . Consequently, instead of imposing random noise on the hidden representation vector, the GASCA model imposes $p_d(x_\mu)$ on $q(y)$ in the following way:

$$q(y) = \mathbb{R}_{x_\mu} q(y|x_\mu) p_d(x_\varphi) dx_\mu \quad (1)$$

With this formulation, the discriminator model D is optimized to rate samples from $p_d(x_\mu)$ with a higher probability, and samples from $q(y)$ with a low probability. Formally this is defined as:

$$\nabla_{\theta_d} \frac{1}{m} \sum_{i=1}^m \left[\log D(x_\mu^{(i)}) + \log (1 - D(G(y^{(i)}))) \right] \quad (2)$$

where x_μ is an input image and x_φ is the target reconstruction image. Note that since $(x_\varphi = x_\mu)$ is not necessarily true, the discriminator D is not guaranteed to see the input to G .

The objective of the autoencoder model G , which in term plays the role as the generator, is to convince the discriminator model D that a sample reconstruction y was drawn from the data distribution $p_d(x_\mu)$ and not from $q(y)$. This optimization is done according to:

$$\nabla_{\theta_g} \frac{1}{m} \sum_{i=1}^m \log(1 - D(G(y^{(i)}))) \quad (3)$$

Furthermore, since GANs are known to be difficult to train due to their sensitivity to hyper-parameters and parameter initialization which often leads to mode collapse, the GASCA model is trained in a GLW fashion. However, since the greedy nature of GLW leads to error accumulation as individual layers are trained and stacked [5], we build on the gradual greedy layer-wise training algorithm from [5] and adapt it for adversarial autoencoders. Accordingly, we introduce the GAN gradual greedy layer-wise (GANGGLW) training framework and formally define it in Algorithm 1.

Algorithm 1 Given a training set X and validation set \tilde{X} each containing input images x_φ and target images x_μ , m shallow autoencoders, an unsupervised feature learning algorithm \mathcal{L} —see Algorithm 2— which returns a trained shallow autoencoder and a discriminator model, and a fine-tuning algorithm \mathcal{T} —see Algorithm 3: train D^1 and G^1 jointly with raw data and add them to their corresponding stacks G and D . For the remaining autoencoders and generator models: encode X and \tilde{X} using the encoder layers ξ from the stack G . Create a new discriminator D^k and train together with the new autoencoder G^k and add them to their corresponding stacks. Fine-tune G on raw pixel data. Forward propagate $x_\varphi \subset X$ through G and use the resulting features, along with x_μ , to fine-tune D for binary classification.

```

1:  $[G^1, D^1] \leftarrow \mathcal{L}(G^1, D^1, X, \tilde{X})$ 
2:  $G \leftarrow G \circ G^1$ 
3:  $D \leftarrow D \circ D^1$ 
4: for  $k = 2, \dots, m$  do
5:    $[\xi, \delta] \leftarrow D$ 
6:    $[X_g, \tilde{X}_g] \leftarrow \xi(X, \tilde{X})$ 
7:    $[G^k, D^k] \leftarrow \mathcal{L}(G^k, D^k, X_d, \tilde{X}_d)$ 
8:    $G \leftarrow G^{(k)} \circ G$ 
9:    $D \leftarrow D^{(k)} \circ D$ 
10:   $G \leftarrow \mathcal{T}(G, X, \tilde{X})$ 
11:   $X_\varphi \leftarrow G(x_\varphi)$ 
12:   $D \leftarrow \mathcal{T}(D, \{X_\varphi, x_\mu \subset X\})$ 
13: end for
14: return  $G, D$ 

```

By fine-tuning G and D in Algorithm 1 we avoid error accumulation from one layer to the next and reduce the required number of fine-tuning steps for deeper layers. However, in

Algorithm 2 Given a training dataset X with m mini-batches of size b , an autoencoder model G and discriminator model D both with weight matrices W_g and W_d , an absolute value cost function $loss$: train G and D jointly such that:

```

1:  $V(W_d) \leftarrow \frac{2}{N_{in} + N_{out}}$ 
2:  $V(W_g) \leftarrow \frac{2}{N_{in} + N_{out}}$ 
3: for  $k = 1, \dots, M$  do
4:   for  $n = 1, \dots, m$  do
5:      $[x_\varphi, x_\mu]_n \subset 1, \dots, m \leftarrow random(X, b)$ 
6:      $y_g \leftarrow predict(x_\varphi, G)$ 
7:      $L_g \leftarrow loss(x_\mu, y_g)$ 
8:      $G \leftarrow update(G, L_g)$ 
9:      $p_\mu \leftarrow predict(x_\mu, D)$ 
10:     $L_{p_d(x_\mu)} \leftarrow loss(1, p_\mu)$ 
11:     $D \leftarrow update(D, L_{p_d(x_\mu)})$ 
12:     $p_y \leftarrow predict(y_g, D)$ 
13:     $L_{q(y)} \leftarrow loss(0, p_y)$ 
14:     $L_{adversary} = L_{p_d(x_\mu)} + L_{q(y)}$ 
15:     $L_{minimax} \leftarrow loss(1, p_y)$ 
16:     $L = L_{minimax} + L_g$ 
17:     $MM_{L_g} \leftarrow lossGrad(1, p_y)$ 
18:     $MM_g \leftarrow Grad(y_g, MM_{L-g}, D)$ 
19:     $G \leftarrow update(G, MM_g)$ 
20:     $Adam(L, G)$ 
21:     $SGD(L_{adversary}, D)$ 
22:   end for
23: end for
24: return  $G, D$ 

```

Algorithm 3 Given a training dataset X with m mini-batches of size b , a validation set \tilde{X} and a model f , train f for M epochs

```

1: for  $k = 1, \dots, M$  do
2:   for  $n = 1, \dots, m$  do
3:      $[x, x_\mu]_n \subset 1, \dots, m \leftarrow random(X, b)$ 
4:      $y \leftarrow predict(x, f)$ 
5:      $L \leftarrow loss(x_\mu, y)$ 
6:      $f \leftarrow update(f, L)$ 
7:   end for
8: end for
9: return  $f$ 

```

the special case where the input and target images are significantly different, GANGGLW is not very compatible with Convolutional Neural Networks (CNNs). CNNs are designed to retain spatial information through filter kernels, whereas in GANGGLW the input and target reconstruction images are not always the same and as such spatial information often needs to be shifted or transformed, or partially ignored. For this reason, and since in our experimental set up we are trying to normalize facial pose, we introduce ConvMLP layers in the next section.




Fig. 2: ConvMPLP layers illustration. Connection weights for the *shifting units* are shared between all the feature maps.

IV. CONVMLP AND HALFCONV LAYERS

One of the main advantages offered by CNNs over multi-layer perceptron networks (MLPs) is their ability to self-learn a translation invariant downsampled feature vector that highlights salient features and retains spatial information through filter kernels. However, CNNs are constrained to preserve the spatial structure of images and therefore are not suitable to reduce or increase facial pose: since every output value produced by convolutional layers is the results of the dot product between a filter kernel and a small view of the input image, the pixel values can only be shifted within the space covered by the filter kernel. Normally, filter kernels tend to be small in order to capture small salient features.

To overcome the limitations imposed by convolutional kernels and fully connected layers, and at the same time exploit the advantages offered by both, we introduce a hybrid layer that combines both approaches. The most straightforward to accomplish this is by simply placing an MLP after the convolutional layer. And, by having a smaller number of hidden units in the MLP than the number of features produced by the convolutional kernels, there would be no need for down-sampling layers such as average or max pooling or convolutional layers with a stride greater than one, which often result in the loss of important information. However, because convolutional layers normally employ a high number of convolutional kernels, this approach would require a significantly large weight matrix W . Accordingly, W would need to have a connection weight for each feature in the feature maps produced by convolutional kernels, resulting in a large number of learnable parameters, increased computational cost, and increased training difficulty.

In contrast, the novel layer presented here, referred to as ConvMPLP hereafter, shapes the resulting feature map produced by a convolution operation with a fully connected layer that is shared between all the resulting feature maps. Refer to Figure 2 for a pictorial description. Given an input image I and a filter kernel K with $m \times n$ dimensions, and a second weight matrix W , the output of ConvMPLP layers is defined as:

$$C(i, j) = W((I * K)(i, j)) \quad (4)$$

where:

$$(I * K)(i, j) = \sum_m \sum_n I(m, n)K(i - m, j - n) \quad (5)$$

Just as in empirical convolutional layers, the non-linearity is provided by a ReLU activation function, extending the above equation to:

$$y = \max(0, C(i, j)) \quad (6)$$

In this formulation of ConvMPLP layers, during the forward pass, the weight matrix W is used to shape every feature map produced by the convolution operation and is updated only once using backpropagation. Sharing this layer across the third dimension—not taking into account the batch dimension for simplicity—its weight matrix is many orders of magnitude smaller than without weight sharing. This also ensures that the *shifting* layer learns to shift all the features highlighted in every feature plane in the same manner. Notice in Figure 2 how the pixels on the second feature map are at a different location.

In addition to ConvMPLP layers, and in order to support the pose invariant training approach and models presented in this work, a second convolutional layer is introduced here. This novel layer, referred to as HalfConv hereafter, exploits facial symmetry present in face images with an estimated pose of zero degrees. HalfConv layers slice the input vector vertically in half. The half containing all the facial features belonging to the left side of a face is then used as input for a convolutional layer that has half the number of parameters than an empirical convolutional layer. The resulting feature map is then simply mirrored across the y axis.

When applied to face or facial expression images, HalfConv layers give up some important information on the right edge of the input image, which in effect corresponds to the features in the middle of a face. This is due to the nature of the convolution operation, which convolves a kernel across an input image, resulting in a feature plane with smaller dimensions than the input image. For this reason, HalfConv layers enforce zero padding p on right side edge of the input image to allow the filter kernel to capture the features closer to the edge. Their output is then defined by:

$$C(i, j) = (I * K)(i, j) = \sum_m \sum_n I(m, n)K(i - m, (j + p) - n) \quad (7)$$

where $p = \frac{j}{2} + 1$. Then every resulting feature plane is reflected over the y axis, resulting in a full image. Note that padding p is enforced to avoid losing features at the edges of the image.

The main advantage offered by HalfConv layers is the reduced number of learnable parameters, which in effect results in easier and faster training. Because the only extra operation required by this layer is simply mirroring a feature vector vertically, HalfConv layers are significantly less computationally expensive than empirical convolutional layers. Furthermore, because this layer only deals with frontal faces, there is no need to employ any shifting neurons. Note that these layers are only suitable for cases where symmetry is existent in the input image or is desired in the resulting feature plane. Therefore, in the GASCA model, these layers are only used when $\alpha = 0$.

V. UNSUPERVISED FEATURE LEARNING

We train two Generative Adversarial Stacked Convolutional Autoencoders, one to normalize facial pose (GASCA₁) and another to normalize facial pose and illumination (GASCA₂). We train both models using the GANGGLW training method



Fig. 3: Top row: input images x_φ to the GASCA₂ model with estimated facial poses at +60, +45, +30, +15, 0 degrees. Bottom row: corresponding reconstructions y produced by the GASCA₂ model with an estimated pose at ~ 0 degrees.

introduced earlier. We train each shallow autoencoder to gradually reduce facial pose, or keep it the same if it is already smaller than the desired target. This process is repeated until reaching a facial pose of 0 degrees. Effectively, the search space for the upper layers is greater than that of the deepest layer, which only has to learn one facial pose of 0 degrees.

For the GASCA₂ model we incorporate illumination invariance normalization training by taking images with disproportionate degrees of illumination and reconstructing them as images with good illumination: good illumination is determined by their relative luminance Y as done by [5].

We employ the MultiPie dataset [22] to train these models given that it contains facial images with multi-pose and multi-illumination.

The training dataset is built according to:

$$x_\mu = \begin{cases} x_{\varphi-d} & , \text{ if } 0 < \alpha < \varphi \\ x_{\varphi+d} & , \text{ if } \varphi < \alpha < 0 \\ x_\varphi & , \text{ if } |\varphi| \leq |\alpha| \end{cases} \quad (8)$$

where $-60 \leq \varphi \leq 60$, α denotes the desired target pose, $\alpha \in \{0, \pm 15, \pm 30, \pm 45\}$, and d denotes the change in pose by degrees: 15 degrees in this work. Each subset is further split into 70% training and 30% validation subsets.

For the GASCA₂ model, since every image in the MultiPie dataset has 19 copies with different levels of relative luminance, we measure their cumulative relative luminance and pick the image closest to the mean as the target reconstruction for all others.

Since $\alpha \in \{0, \pm 15, \pm 30, \pm 45\}$, we design the GASCA models with three ConvMLP layers and one HalfConv Layer for the encoder element. ConvMLP layers use 5×5 , 3×3 , 3×3 filter kernels and 100 hidden units for the *shifting* units. The decoder element only uses deconvolutional layers to force the encoder to learn a downsampled pose invariant hidden representation, since we fine-tune it later on to do classification. Accordingly, every shallow autoencoder is trained on a single target pose. For instance, subset A_1 contains all the images with $\{0, \pm 15\}$ degrees. A_2 contains all images at angles $\{0, \pm 15, \pm 30\}$, thus $A_1 \cap A_2$,

Every shallow autoencoder in the GASCA models is trained for 100 and fine-tuned for 20 epochs. G is optimized using ADAM [23], whereas D employs SGD with Nesterov momentum. The initial learning rates for each individual shallow autoencoder in G were set to $\lambda \in \{0.1, 0.3, 0.5, 0.7, 0.75\}$. Since D learns faster than G , the shallow autoencoders employ smaller learning rates: $\lambda \in \{0.01, 0.03, 0.5, 0.07\}$. During fine-tuning, the stacks G and D use a learning rate of 0.001. This combination of hyper-parameters provided the best results for both models.

VI. POSE INVARIANT RECONSTRUCTION RESULTS

The novel pose invariant Generative Adversarial Stacked Convolutional Autoencoder models proposed in this work are trained to gradually reduce facial pose using GANGGLW training. As it can be observed in Figure 3, the pose and illumination invariant GASCA₂ model manages to reduce facial pose in facial images with an estimated pose of up to ± 60 degrees. It also produces reconstructions with similar illumination.

It can also be observed that on the images with pose of ± 60 degrees half of the face is not visible, yet the pose invariant model manages to fill in the missing information, and more importantly keeps the shape of facial shapes which are important for emotion recognition: eyes, eyebrows, mouth, nose, cheeks, among others. Nonetheless, the greater the pose in x_φ the poorer the quality of the reconstruction y . This is justified by (i) the fact that the model has to compensate for missing information, (ii) the fact that only one layer is trained specifically to deal with that particular facial pose, (iii) the smaller the pose the more the images get seen by every layer in G during training, and (iv) increased network depth.

If the shallow autoencoder at step $k = 1$ fails to learn a pose invariant feature vector, the shallow autoencoder at step $k = 2$ will struggle even more to learn a pose invariant feature vector, and so forth. GANGGLW greatly helps to address this issue by allowing inter-layer fine-tuning, which helps strengthen the weight connections between D^k and D^{k+1} .

One of the main remarks observed in the reconstructions is that although these retain all the important salient features, they are visually different than the input images. These reconstructions could be improved by unsupervised fine-tuning of G for a significantly longer number of epochs. Likewise, secondary methods such as super resolution CNNs [24] could be used to improve the visual quality of the reconstructed images. However, because the objective of this research is to only learn a pose invariant feature vector z that can be used for emotion recognition, the quality or resolution of the reconstructions is trivial.

One of the main advantages offered by ConvMLP layers is that the number of *shifting* neurons can be adjusted as needed. In the GASCA models, every ConvMLP layer only employs 100, which are enough to reposition facial features and eventually reduce facial pose. Another advantage offered by ConvMLP layers is that they can be used for dimensionality

reduction by mapping a feature plane to a smaller feature plane. Although, this is not evaluated in this research.

As illustrated in Figure 3, the reconstructed images also do not have a horizontal line diving the face in two, as it would be expected due to the use of HalfConv layers. When visualizing the feature planes produced by these layers, the line is somewhat visible. However, because in the final stack G this layer is followed by all the layers in the decoder stack of G , and since the line is not visible in the target reconstruction images, it vanishes during fine-tuning.

We did not notice significant differences between the reconstructions of both GASCA₁ and GASCA₂ models. However, the reconstruction loss for the latter was marginally smaller and as seen in the next section it generalizes better.

VII. POSE AND ILLUMINATION INVARIANT EMOTION RECOGNITION

Once a GASCA model is trained and fine-tuned for reconstruction, it can be used as a generic feature extractor for facial expression images. However, it can only provide feature vectors that are pose and illumination invariant but that do not necessarily discriminate between different emotions. Therefore, we fine-tune the encoder element of G for classification. We discard both the discriminator model D along with the decoder element g_D of the generator D and attach a classification layer to the encoder.

The pose invariant GASCA₁ model is used to initialize a classifier model, CNN₁, which is fine-tuned and tested on the KDEF corpus [25]. This dataset contains frontal and images at ± 45 degrees. No other publicly available datasets with multiple poses have facial expression labels.

The pose and illumination invariant GASCA₂ model is used to initialize a second classifier, CNN_{2a}. This model is also fine-tuned and tested on the KDEF corpus.

In addition, in an attempt to test the robustness of our proposed methodology, we use GASCA₂ to initialize a third model CNN_{2b}. However, due to the lack of publicly available data taken in realistic environments with multi-pose and varying illumination, as well as labels for the emotions being expressed, we build a large dataset composed of the CK+ [26], JAFFE [27], KDEF [25], and FEEDTUM [28] corpora. We refer to this corpus as combined facial expressions (CFE). Note that because, as later discussed, we obtain over 99.6% on this corpus, this model is evaluated on completely novel data: the entire NAOFaces corpus [29]. This set has total of 196 images collected in unconstrained environments. Participants were 28 21 males and 7 females between ages 18 and 55 from at least five different ethnic backgrounds.

As opposed to empirical CNN classifier models which employ a fully connected layer after the last convolutional layer, the classifiers in this work map the resulting feature planes produced by the last convolutional layer, which is a HalfConv layer, directly to an output SoftMax layer for classification, as done in [30].

The CNN₁ and CNN_{2a} models are fine-tuned for 10 epochs and, because the CFE corpus has more images, CNN_{2b} is only fine-tuned for two epochs. Since the stacked autoencoders are optimized using ADAM, all classifiers are fine-tuned also using ADAM and a learning rate of 0.01. Using a different optimizer like SGD for fine-tuning would lead to the gradients changing drastically and require a longer fine-tuning process.

VIII. POSE INVARIANT EMOTION RECOGNITION RESULTS

As it can be observed in Table I, the pose invariant model, CNN₁, obtains a classification performance of 96.81%. In contrast, the pose invariant model that also incorporated illumination invariance obtains a state-of-the-art classification rate of 98.07%. The main differences in performance are observed for classes: surprise, Fear, and Angry, whereas both CNN₁ and CNN_{2a} obtained the same classification accuracy for the remaining classes. Because both models are trained using a relatively similar approach, it is hypothesized that these discrepancies in classification performance are due to these three classes containing more images with varying image luminance, thus the pose and illumination invariant model is able to generalize better.

The CNN_{2b} model is evaluated on the NAOFaces corpus and achieves 81.36% accuracy. This is significantly lower than the performance of the other models on the KDEF corpus. We attribute this lower performance to the fact that the NAOFaces contains images that are substantially more difficult, i.e. people with glasses, at different poses, and different ethnicity. Moreover, this model was not fine-tuned on any images from this corpus. This theory is further supported by the fact that when we split the CFE corpus 80% training and 20% testing, we obtain 99% on the test set.

One important observation in Table II is that, when looking at the missclassified images for a given class, on average 40% of them are frontal images, i.e. images with zero degrees pose, and the remaining 60% are those with a pose. However, because the ratio of images with a facial pose is 2:1 compared to those without one. This means that on average, more images without facial pose are missclassified. These results and observations are of great importance given that they support the pose invariant pretraining approach presented in this work. Another observation is that not a single image from the other classes was confused with Neutral. This particular score is significant taking into account that all emotions derive from a neutral state, often resulting in low precision scores.

Despite the good performance offered by the CNN_{2a} on the NAOFaces corpus, the classification performance offered by this model is not ideal. This is attributed to one major factor: cultural differences. Because the model was trained solely on images from Caucasian people, the model has never learned to adjust to cultural difference. The NAOFaces corpus contains images of people from at least five different backgrounds including: Asian, Arab, Black, Irish and Hispanic, among others unrevealed ones. In effect, because people from different ethnic backgrounds express emotions differently [31],

TABLE I: (left) Classification performance (96.810%) of the CNN_1 model on the KDEF corpus pretrained for pose invariance. (right) Classification performance (98.070%) of the CNN_{2a} model on the KDEF corpus pretrained for pose and illumination invariance.

	A	D	F	H	N	Sa	Su		A	D	F	H	N	Sa	Su
A	94.44	1.59	1.59	0.00	0.79	1.59	0.00		96.83	0.79	1.59	0	0.00	0.79	0.00
D	0.00	97.60	0.00	0.00	0.00	2.40	0.00		0.00	97.60	0.00	0.00	0.00	2.40	0.00
F	0.00	0.79	89.68	0.79	0.00	3.97	4.76		0.00	0.79	93.65	0.79	0.00	2.38	2.38
H	0.00	0.00	0.00	100.00	0.00	0.00	0.00		0.00	0.00	0.00	100.00	0.00	0.00	0.00
N	0.00	0.00	0.00	0.00	100.00	0.00	0.00		0.00	0.00	0.00	0.00	100.00	0.00	0.00
Sa	0.79	0.79	0.00	0.00	0.00	98.41	0.00		0.79	0.79	0.00	0.00	0.00	98.41	0.00
Su	0.00	0.00	2.42	0.00	0.00	0.00	97.58		0.00	0.00	0.00	0.00	0.00	0.00	100.00

TABLE II: Classification performance (81.36%) of the CNN_{2b} model on the NAOFaces corpus.

	A	D	F	H	N	Sa	Su
A	92.86	7.14	0.00	0.00	0.00	0.00	0.00
D	8.33	75.00	8.33	0.00	0.00	8.33	0.00
F	9.09	0.00	81.81	0.00	0.00	0.00	9.09
H	0.00	0.00	0.00	100.00	0.00	0.00	0.00
N	3.85	0.00	3.85	15.38	57.69	11.54	7.69
Sa	9.09	0.00	18.18	0.00	0.00	72.72	0.00
Su	0.00	0.00	10.53	0.00	0.00	0.00	89.47

TABLE III: Classification performance comparison on the KDEF corpus: ResNet-34 —state-of-the-art classifier; CNN_1 —pose invariant classifier proposed; CNN_{2a} pose and illumination invariant classifier proposed.

	<i>Resnet34</i>	CNN_1	CNN_{2a}
A	84.127%	94.444%	96.825%
D	85.600%	97.600%	97.600%
F	73.810%	89.683%	93.651%
H	98.413%	100.000%	100.000%
N	90.400%	100.000%	100.000%
Sa	84.921%	98.413%	98.413%
Su	95.161%	97.581%	100.000%
<i>Total</i>	87.472%	96.810%	98.070%

the classifier should be trained with images of participants from a wide range of ethnic backgrounds and cultures.

IX. COMPARISON AGAINST STATE-OF-THE-ART

We now compare our methodology to contemporary state-of-the-art methods on the KDEF corpus. Due to the lack of contemporary work designed explicitly for pose invariant emotion recognition, the methods proposed in this work are compared against one of the most common and state-of-the-art classifiers: a ResNet [30]. Accordingly, a ResNet-34, i.e. with 34 parametrised layers, is trained using SGD, a momentum of 0.9 and learning rate of 0.1. This model is trained for 100 on the training subset of the KDEF corpus and achieves an accuracy rate of 87.472% on the test subset, as illustrated in Table III. Note that even though the authors of [8] report 92.52% on the KDEF corpus, those results are only reported on frontal faces without facial pose. On the contrary, all the models in this section are evaluated on images with multiple poses, hence the marginally lower performance of the ResNet model.

As seen in Table III, the pose and illumination invariance model, CNN_{2a} outperforms the state-of-the-art classifier ResNet-34 model by over 10%. Similarly, it outperforms CNN_1 marginally, supporting the pose and illumination invariant training approach. The pose invariant GASCA models also have an exponentially smaller number of parameters compared to the ResNet-34 model.

The novelty of this work also arises from combining greedy layer-wise training with adversarial learning. Generative Adversarial Autoencoders are trained jointly as opposed to layer-wise. They impose a random distribution $p(z)$ on the distribution $q(z)$ produced by the encoder element of G , and use the resulting aggregate posterior distribution is

mapped to reconstruction y . The discriminator D tries to guess if the sample was drawn from $q(z)$ or $p(z)$. The GASCA models do not use a random distribution and instead use the reconstruction y produced by forward propagating x_φ through G , along with the target image x_μ as input for the discriminator. The generator G is optimized to reduce the distance between y and x_μ . By fine-tuning the stacks G and D at every step k , both models become better at their respective job. By improving the ability of D to differentiate between y and x_μ , G is forced to produce remarkable reconstructions and learn an encoder function that produces downsampled pose invariant feature vectors.

In terms of work on pose reductions, a similar model was proposed by [32]. However, the authors focused on face detection and their model does not make use of Convolutional Autoencoders and instead uses MLPs, which are prone to overfitting when applied to this problem. Furthermore, because their model does not take into account spatial information, it is unable to retain salient features that are essential for emotion recognition. Whereas the GASCA models are able to retain facial features, or compensate for missing information when this is not present in the image. Additionally, the $GASCA_{2a}$ model also takes into account illumination and produces an illumination and pose invariant feature vector.

X. CONCLUSIONS AND FUTURE DIRECTIONS

This work has introduced a novel pose and illumination invariant facial expression recognition model. A CNN classifier is pretrained as a Generative Adversarial Stacked Convolutional Autoencoder in a gradual greedy layer-wise semi-supervised fashion. The GASCA model learns to map an input

image containing a face, with an estimate pose φ , to a hidden representation z with an estimated pose of 0 degrees. Once the GASCA model is trained, the encoder element is used to initialize a CNN model which is fine-tuned for classification.

The outstanding performance of the GASCA models is derived from four concepts: (i) our GANGGLW training method (ii) the ConvMLP layers with *shifting* neurons, (iii) the HalfConv layers which take exploit of facial symmetry, and (iv) multi-pose facial expressions data. Our pose and illumination invariant method produces state-of-the-art classification performance on multi-pose facial expression corpora. Moreover, the GASCA model produces reconstruction with very small errors and is able to generalize on unseen data.

The success of the pose invariant models is in part due to ConvMLP layers, which learn salient features and shift them as needed to reduce facial pose. HalfConv layers also play an important role as they reduce the number of learning parameters. HalfConv layers were inspired by the model presented by [33], which splits the input images in half to simplify feature learning.

To the best of the authors' knowledge, this is the first approach that combines a greedy layer-wise training method with adversarial learning. This is also the first approach to solely focus on pose and illumination invariant emotion recognition. Future work will look at exploiting the ability of our model to generate new data in order to deal with scenarios where lack of multi-pose labeled exists.

LIST OF REFERENCES

- [1] G. Levi and T. Hassner, "Emotion Recognition in the Wild via Convolutional Neural Networks and Mapped Binary Patterns," in *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction - ICMI '15*, (New York, New York, USA), pp. 503–510, ACM Press, 2015.
- [2] A. Krizhevsky, L. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *NIPS*, pp. 1106–1114, 2012.
- [3] S. Ioffe and C. Szegedy, "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift," feb 2015.
- [4] P. Vincent PASCALVINCENT and H. Larochelle LAROCHEH, "Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion Pierre-Antoine Manzagol," *Journal of Machine Learning Research*, vol. 11, pp. 3371–3408, 2010.
- [5] A. Ruiz-Garcia, V. Palade, M. Elshaw, and I. Almakky, "Deep Learning for Illumination Invariant Facial Expression Recognition," in *Proceedings of the International Joint Conference on Neural Networks*, (Rio de Janeiro), IEEE, 2018.
- [6] I. Goodfellow, Bengio, Yoshua, and A. Courville, *Deep Learning*. MIT Press, 2016.
- [7] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle, "Greedy Layer-Wise Training of Deep Networks," tech. rep.
- [8] A. Ruiz-Garcia, M. Elshaw, A. Altahhan, and V. Palade, "Stacked deep convolutional auto-encoders for emotion recognition from facial expressions," in *Proceedings of the International Joint Conference on Neural Networks*, vol. 2017-May, pp. 1586–1593, IEEE, may 2017.
- [9] A. Makhzani, J. Shlens, N. Jaitly, I. Goodfellow, and B. Frey, "Adversarial Autoencoders," 2015.
- [10] I. Goodfellow, J. Pouget-Abadie, M. M. A. in neural ..., and U. 2014, "Generative adversarial nets," *Advances in Neural Information Processing Systems* 27, pp. 2672–2680, 2014.
- [11] S. A. Israel, J. Goldstein, J. S. Klein, J. Talamonti, F. Tanner, S. Zabel, P. A. Sallee, and L. McCoy, "Generative Adversarial Networks for Classification," in *2017 IEEE Applied Imagery Pattern Recognition Workshop (AIPR)*, pp. 1–4, IEEE, oct 2017.
- [12] A. Chatziagapi, G. Paraskevopoulos, D. Sgouropoulos, G. Pantazopoulos, M. Nikandrou, T. Giannakopoulos, A. Katsamanis, A. Potamianos, and S. Narayanan, "Data Augmentation using GANs for Speech Emotion Recognition," 2019.
- [13] W. Yi, Y. Sun, and S. He, "Data Augmentation Using Conditional GANs for Facial Emotion Recognition," in *Progress in Electromagnetics Research Symposium*, vol. 2018-Augus, pp. 710–714, Institute of Electrical and Electronics Engineers Inc., dec 2018.
- [14] X. Zhu, Y. Liu, Z. Qin, and J. Li, "Data Augmentation in Emotion Classification Using Generative Adversarial Networks," nov 2017.
- [15] L. Tran, X. Yin, and X. Liu, "Disentangled representation learning GAN for pose-invariant face recognition," in *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, vol. 2017-Janua, pp. 1283–1292, 2017.
- [16] J. Zhao, L. Xiong, K. Jayashree, J. Li, F. Zhao, Z. Wang, S. Pranata, S. Shen, S. Yan, and J. Feng, "Dual-Agent GANs for Photorealistic and Identity Preserving Profile Face Synthesis," *Nips 2017*, no. 15, pp. 1–11, 2017.
- [17] J. Chen, J. Konrad, and P. Ishwar, "VGAN-Based Image Representation Learning for Privacy-Preserving Facial Expression Recognition," tech. rep.
- [18] D. H. Liu, K. M. Lam, and L. S. Shen, "Illumination invariant face recognition," *Pattern Recognition*, vol. 38, pp. 1705–1716, oct 2005.
- [19] C. Tosik, A. Eleyan, and M. Salman, "Illumination invariant face recognition system," in *2013 21st Signal Processing and Communications Applications Conference, SIU 2013*, pp. 1–4, IEEE, apr 2013.
- [20] O. Gupta, D. Raviv, and R. Raskar, "Deep video gesture recognition using illumination invariants," *Arxiv*, pp. 1–9, 2016.
- [21] X. Chen, X. Lan, G. Liang, J. Liu, and N. Zheng, "Pose-and-illumination-invariant face representation via a triplet-loss trained deep reconstruction model," *Multimedia Tools and Applications*, vol. 76, pp. 22043–22058, nov 2017.
- [22] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker, "Multi-PIE," in *2008 8th IEEE International Conference on Automatic Face & Gesture Recognition*, pp. 1–8, IEEE, sep 2008.
- [23] D. P. Kingma and J. L. Ba, "Adam," pp. 1–15, 2015.
- [24] W. Shi, J. Caballero, F. Huszar, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, "Real-Time Single Image and Video Super-Resolution Using an Efficient Sub-Pixel Convolutional Neural Network," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1874–1883, 2016.
- [25] D. Lundqvist, A. Flykt, and A. Öhman, "The Karolinska Directed Emotional Faces - KDEF CD ROM from Department of Clinical Neuroscience, Psychology section," *Karolinska Institutet*, pp. 3–5, 1998.
- [26] P. Lucey, J. F. Cohn, T. Kanade, J. Saragih, Z. Ambadar, and I. Matthews, "The Extended Cohn-Kanade Dataset (CK+): A complete dataset for action unit and emotion-specified expression," tech. rep.
- [27] M. Lyons, S. Akamatsu, M. Kamachi, and J. Gyoba, "Coding facial expressions with Gabor wavelets," in *Proceedings - 3rd IEEE International Conference on Automatic Face and Gesture Recognition, FG 1998*, pp. 200–205, 1998.
- [28] F. Wallhoff, B. Schuller, M. Hawellek, and G. Rigoll, "Efficient recognition of authentic dynamic facial expressions on the feedtum database," in *2006 IEEE International Conference on Multimedia and Expo, ICME 2006 - Proceedings*, vol. 2006, pp. 493–496, IEEE, jul 2006.
- [29] A. Ruiz-Garcia, N. Webb, V. Palade, M. Eastwood, and M. Elshaw, "Deep Learning for Real Time Facial Expression Recognition in Social Robots," in *Proceedings of the International Conference on Neural Information Processing*, 2018.
- [30] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," *Arxiv.Org*, vol. 7, pp. 171–180, dec 2015.
- [31] N. M. Hewahi and A. R. M. Baraka, "Impact of Ethnic Group on Human Emotion Recognition Using Backpropagation Neural Network," *BRAIN. Broad Research in Artificial*, pp. 20–27, 2012.
- [32] M. Kan, S. Shan, H. Chang, and X. Chen, "Stacked progressive auto-encoders (SPAe) for face recognition across poses," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 1883–1890, IEEE, jun 2014.
- [33] A. Ruiz-Garcia, M. Elshaw, A. Altahhan, and V. Palade, "Emotion Recognition Using Facial Expression Images for a Robotic Companion," in *Engineering Applications of Neural Networks: 17th International Conference, EANN 2016, Aberdeen, UK, September 2-5, 2016, Proceedings*, pp. 79–93, Springer, Cham, 2016.