

GANimation: Anatomically-aware Facial Animation from a Single Image

Albert Pumarola¹, Antonio Agudo¹, Aleix M. Martinez²,
Alberto Sanfeliu¹, Francesc Moreno-Noguer¹

¹Institut de Robòtica i Informàtica Industrial, CSIC-UPC, 08028, Barcelona, Spain

²The Ohio State University, Columbus, OH 43210, USA

Abstract. Recent advances in Generative Adversarial Networks (GANs) have shown impressive results for task of facial expression synthesis. The most successful architecture is StarGAN [4], that conditions GANs' generation process with images of a specific domain, namely a set of images of persons sharing the same expression. While effective, this approach can only generate a discrete number of expressions, determined by the content of the dataset. To address this limitation, in this paper, we introduce a novel GAN conditioning scheme based on Action Units (AU) annotations, which describes in a continuous manifold the anatomical facial movements defining a human expression. Our approach allows controlling the magnitude of activation of each AU and combine several of them. Additionally, we propose a fully unsupervised strategy to train the model, that only requires images annotated with their activated AUs, and exploit attention mechanisms that make our network robust to changing backgrounds and lighting conditions. Extensive evaluation show that our approach goes beyond competing conditional generators both in the capability to synthesize a much wider range of expressions ruled by anatomically feasible muscle movements, as in the capacity of dealing with images in the wild.

Keywords: GANs, Face Animation, Action-Unit Condition.

1 Introduction

Being able to automatically animate the facial expression from a single image would open the door to many new exciting applications in different areas, including the movie industry, photography technologies, fashion and e-commerce business, to name but a few. As Generative and Adversarial Networks have become more prevalent, this task has experienced significant advances, with architectures such as StarGAN [4], which is able not only to synthesize novel expressions, but also to change other attributes of the face, such as age, hair color or gender. Despite its generality, StarGAN can only change a particular aspect of a face among a discrete number of attributes defined by the annotation granularity of the dataset. For instance, for the facial expression synthesis task,



Fig. 1. Facial animation from a single image. We propose an anatomically coherent approach that is not constrained to a discrete number of expressions and can animate a given image and render novel expressions in a continuum. In these examples, we are given solely the left-most input image \mathbf{I}_{y_r} (highlighted by a green square), and the parameter α controls the degree of activation of the target action units involved in a smiling-like expression. Additionally, our system can handle images with unnatural illumination conditions, such as the example in the bottom row.

[4] is trained on the RaFD [16] dataset which has only 8 binary labels for facial expressions, namely sad, neutral, angry, contemptuous, disgusted, surprised, fearful and happy.

Facial expressions, however, are the result of the combined and coordinated action of facial muscles that cannot be categorized in a discrete and low number of classes. Ekman and Friesen [6] developed the Facial Action Coding System (FACS) for describing facial expressions in terms of the so-called Action Units (AUs), which are anatomically related to the contractions of specific facial muscles. Although the number of action units is relatively small (30 AUs were found to be anatomically related to the contraction of specific facial muscles), more than 7,000 different AU combinations have been observed [30]. For example, the facial expression for *fear* is generally produced with activations: Inner Brow Raiser (AU1), Outer Brow Raiser (AU2), Brow Lowerer (AU4), Upper Lid Raiser (AU5), Lid Tightener (AU7), Lip Stretcher (AU20) and Jaw Drop (AU26) [5]. Depending on the magnitude of each AU, the expression will transmit the emotion of fear to a greater or lesser extent.

In this paper we aim at building a model for synthetic facial animation with the level of expressiveness of FACS, and being able to generate anatomically-

aware expressions in a continuous domain, without the need of obtaining any facial landmarks [36]. For this purpose we leverage on the recent EmotioNet dataset [3], which consists of one million images of facial expressions (we use 200,000 of them) of emotion in the wild annotated with discrete AUs activations¹. We build a GAN architecture which, instead of being conditioned with images of a specific domain as in [4], it is conditioned on a one-dimensional vector indicating the presence/absence and the magnitude of each action unit. We train this architecture in an unsupervised manner that only requires images with their activated AUs. To circumvent the need for pairs of training images of the same person under different expressions, we split the problem in two main stages. First, we consider an AU-conditioned bidirectional adversarial architecture which, given a single training photo, initially renders a new image under the desired expression. This synthesized image is then rendered-back to the original pose, hence being directly comparable to the input image. We incorporate very recent losses to assess the photorealism of the generated image. Additionally, our system also goes beyond state-of-the-art in that it can handle images under changing backgrounds and illumination conditions. We achieve this by means of an attention layer that focuses the action of the network only in those regions of the image that are relevant to convey the novel expression.

As a result, we build an anatomically coherent facial expression synthesis method, able to render images in a continuous domain, and which can handle images in the wild with complex backgrounds and illumination conditions. As we will show in the results section, it compares favorably to other conditioned-GANs schemes, both in terms of the visual quality of the results, and the possibilities of generation. Figure 1 shows some example of the results we obtain, in which given one input image, we gradually change the magnitude of activation of the AUs used to produce a smile.

2 Related Work

Generative Adversarial Networks. GANs are a powerful class of generative models based on game theory. A typical GAN optimization consists in simultaneously training a generator network to produce realistic fake samples and a discriminator network trained to distinguish between real and fake data. This idea is embedded by the so-called *adversarial loss*. Recent works [1,9] have shown improved stability relying on the continuous Earth Mover Distance metric, which we shall use in this paper to train our model. GANs have been shown to produce very realistic images with a high level of detail and have been successfully used for image translation [38,10,13], face generation [12,28], super-resolution imaging [34,18], indoor scene modeling [12,33] and human poses editing [27].

Conditional GANs. An active area of research is designing GAN models that incorporate conditions and constraints into the generation process. Prior studies have explored combining several conditions, such as text descriptions [29,39,37]

¹ The dataset was re-annotated with [2] to obtain continuous activation annotations.

and class information [24,23]. Particularly interesting for this work are those methods exploring image based conditioning as in image super-resolution [18], future frame prediction [22], image in-painting [25], image-to-image translation [10] and multi-target domain transfer [4].

Unpaired Image-to-Image Translation. As in our framework, several works have also tackled the problem of using unpaired training data. First attempts [21] relied on Markov random field priors for Bayesian based generation models using images from the marginal distributions in individual domains. Others explored enhancing GANs with Variational Auto-Encoder strategies [21,15]. Later, several works [25,19] have exploited the idea of driving the system to produce mappings transforming the style without altering the original input image content. Our approach is more related to those works exploiting cycle consistency to preserve key attributes between the input and the mapped image, such as CycleGAN [38], DiscoGAN [13] and StarGAN [4].

Face Image Manipulation. Face generation and editing is a well-studied topic in computer vision and generative models. Most works have tackled the task on attribute editing [17,26,31] trying to modify attribute categories such as adding glasses, changing color hair, gender swapping and aging. The works that are most related to ours are those synthesizing facial expressions. Early approaches addressed the problem using mass-and-spring models to physically approximate skin and muscle movement [7]. The problem with this approach is that is difficult to generate natural looking facial expressions as there are many subtle skin movements that are difficult to render with simple spring models. Another line of research relied on 2D and 3D morphings [35], but produced strong artifacts around the region boundaries and was not able to model illumination changes.

More recent works [4,24,20] train highly complex convolutional networks able to work with images in the wild. However, these approaches have been conditioned on discrete emotion categories (e.g., happy, neutral, and sad). Instead, our model resumes the idea of modeling skin and muscles, but we integrate it in modern deep learning machinery. More specifically, we learn a GAN model conditioned on a continuous embedding of muscle movements, allowing to generate a large range of anatomically possible face expressions as well as smooth facial movement transitions in video sequences.

3 Problem Formulation

Let us define an input RGB image as $\mathbf{I}_{\mathbf{y}_r} \in \mathbb{R}^{H \times W \times 3}$, captured under an arbitrary facial expression. Every gesture expression is encoded by means of a set of N action units $\mathbf{y}_r = (y_1, \dots, y_N)^\top$, where each y_n denotes a normalized value between 0 and 1 to modulate the magnitude of the n -th action unit. It is worth pointing out that thanks to this continuous representation, a natural interpolation can be done between different expressions, allowing to render a wide range of realistic and smooth facial expressions.

Our aim is to learn a mapping \mathcal{M} to translate $\mathbf{I}_{\mathbf{y}_r}$ into an output image $\mathbf{I}_{\mathbf{y}_g}$ conditioned on an action-unit target \mathbf{y}_g , i.e., we seek to estimate the mapping

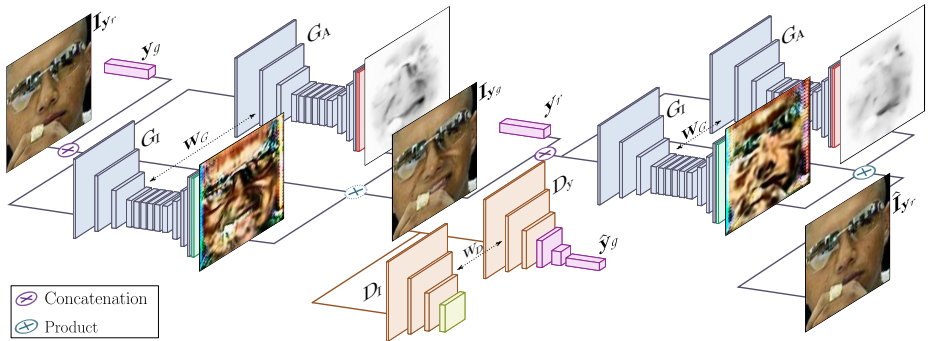


Fig. 2. Overview of our approach to generate photo-realistic conditioned images. The proposed architecture consists of two main blocks: a generator G to regress attention and color masks; and a critic D to evaluate the generated image in its photorealism D_I and expression conditioning fulfillment $\hat{\mathbf{y}}_g$. Note that our systems does not require supervision, i.e., no pairs of images of the same person with different expressions, nor the target image $\mathbf{I}_{\mathbf{y}_g}$ are assumed to be known.

$\mathcal{M} : (\mathbf{I}_{\mathbf{y}_r}, \mathbf{y}_g) \rightarrow \mathbf{I}_{\mathbf{y}_g}$. To this end, we propose to train \mathcal{M} in an unsupervised manner, using M training triplets $\{\mathbf{I}_{\mathbf{y}_r}^m, \mathbf{y}_r^m, \mathbf{y}_g^m\}_{m=1}^M$, where the target vectors \mathbf{y}_g^m are randomly generated. Importantly, we neither require pairs of images of the same person under different expressions, nor the expected target image $\mathbf{I}_{\mathbf{y}_g}$.

4 Our Approach

This section describes our novel approach to generate photo-realistic conditioned images, which, as shown in Fig. 2, consists of two main modules. On the one hand, a generator $G(\mathbf{I}_{\mathbf{y}_r} | \mathbf{y}_g)$ is trained to realistically transform the facial expression in image $\mathbf{I}_{\mathbf{y}_r}$ to the desired \mathbf{y}_g . Note that G is applied twice, first to map the input image $\mathbf{I}_{\mathbf{y}_r} \rightarrow \mathbf{I}_{\mathbf{y}_g}$, and then to render it back $\mathbf{I}_{\mathbf{y}_g} \rightarrow \hat{\mathbf{I}}_{\mathbf{y}_r}$. On the other hand, we use a WGAN-GP [9] based critic $D(\mathbf{I}_{\mathbf{y}_g})$ to evaluate the quality of the generated image as well as its expression.

4.1 Network Architecture

Generator. Let G be the generator block. Since it will be applied bidirectionally (i.e., to map either input image to desired expression and vice-versa) in the following discussion we use subscripts o and f to indicate *origin* and *final*.

Given the image $\mathbf{I}_{\mathbf{y}_o} \in \mathbb{R}^{H \times W \times 3}$ and the N -vector \mathbf{y}_f encoding the desired expression, we form the input of generator as a concatenation $(\mathbf{I}_{\mathbf{y}_o}, \mathbf{y}_o) \in \mathbb{R}^{H \times W \times (N+3)}$, where \mathbf{y}_o has been represented as N arrays of size $H \times W$.

One key ingredient of our system is to make G focus only on those regions of the image that are responsible of synthesizing the novel expression and keep the rest elements of the image such as hair, glasses, hats or jewelry untouched. For this purpose, we have embedded an attention mechanism to the generator.

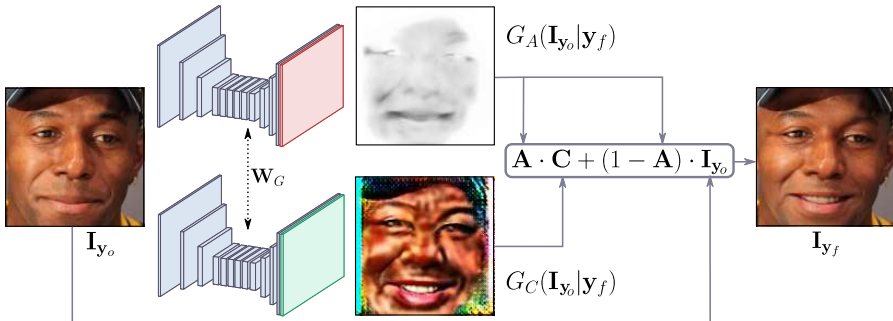


Fig. 3. Attention-based generator. Given an input image and the target expression, the generator regresses an attention mask \mathbf{A} and an RGB color transformation \mathbf{C} over the entire image. The attention mask defines a per pixel intensity specifying to which extend each pixel of the original image will contribute in the final rendered image.

Concretely, instead of regressing a full image, our generator outputs two masks, a color mask \mathbf{C} and attention mask \mathbf{A} . The final image can be obtained as:

$$\mathbf{I}_{y_f} = \mathbf{A} \cdot \mathbf{C} + (1 - \mathbf{A}) \cdot \mathbf{I}_{y_o}, \quad (1)$$

where $\mathbf{A} = G_A(\mathbf{I}_{y_o}|y_f) \in \{0, \dots, 1\}^{H \times W}$ and $\mathbf{C} = G_C(\mathbf{I}_{y_o}|y_f) \in \mathbb{R}^{H \times W \times 3}$. The mask \mathbf{A} indicates to which extend each pixel of the \mathbf{C} contributes to the output image \mathbf{I}_{y_f} . In this way, the generator does not need to render static elements, and can focus exclusively on the pixels defining the facial movements, leading to sharper and more realistic synthetic images. This process is depicted in Fig. 3.

Conditional Critic. This is a network trained to evaluate the generated images in terms of their photo-realism and desired expression fulfillment. The structure of $D(\mathbf{I})$ resembles that of the PatchGAN [10] network mapping from the input image \mathbf{I} to a matrix $\mathbf{Y}_1 \in \mathbb{R}^{H/2^6 \times W/2^6}$, where $\mathbf{Y}_1[i, j]$ represents the probability of the overlapping patch ij to be real. Also, to evaluate its conditioning, on top of it we add an auxiliary regression head that estimates the AUs activations $\hat{\mathbf{y}} = (\hat{y}_1, \dots, \hat{y}_N)^\top$ in the image.

4.2 Learning the Model

The loss function we define contains four terms, namely an *image adversarial loss* [1] with the modification proposed by Gulrajani *et al.* [9] that pushes the distribution of the generated images to the distribution of the training images; the *attention loss* to drive the attention masks to be smooth and prevent them from saturating; the *conditional expression loss* that conditions the expression of the generated images to be similar to the desired one; and the *identity loss* that favors to preserve the person texture identity.

Image Adversarial Loss. In order to learn the parameters of the generator G , we use the modification of the standard GAN algorithm [8] proposed by WGAN-GP [9]. Specifically, the original GAN formulation is based on the

Jensen-Shannon (JS) divergence loss function and aims to maximize the probability of correctly classifying real and rendered images while the generator tries to foul the discriminator. This loss is potentially not continuous with respect to the generators parameters and can locally saturate leading to vanishing gradients in the discriminator. This is addressed in WGAN [1] by replacing JS with the continuous Earth Mover Distance. To maintain a Lipschitz constraint, WGAN-GP [9] proposes to add a gradient penalty for the critic network computed as the norm of the gradients with respect to the critic input.

Formally, let \mathbf{I}_{y_o} be the input image with the initial condition \mathbf{y}_o , \mathbf{y}_f the desired final condition, \mathbb{P}_o the data distribution of the input image, and $\mathbb{P}_{\tilde{T}}$ the random interpolation distribution. Then, the *critic loss* $\mathcal{L}_I(G, D_I, \mathbf{I}_{y_o}, \mathbf{y}_f)$ we use is:

$$\mathbb{E}_{\mathbf{I}_{y_o} \sim \mathbb{P}_o} [D_I(G(\mathbf{I}_{y_o} | \mathbf{y}_f))] - \mathbb{E}_{\mathbf{I}_{y_o} \sim \mathbb{P}_o} [D_I(\mathbf{I}_{y_o})] + \lambda_{\text{gp}} \mathbb{E}_{\tilde{T} \sim \mathbb{P}_{\tilde{T}}} \left[(\|\nabla_{\tilde{T}} D_I(\tilde{T})\|_2 - 1)^2 \right],$$

where λ_{gp} is a penalty coefficient.

Attention Loss. When training the model we do not have ground-truth annotation for the attention masks \mathbf{A} . Similarly as for the color masks \mathbf{C} , they are learned from the resulting gradients of the critic module and the rest of the losses. However, the attention masks can easily saturate to 1 which makes that $\mathbf{I}_{y_o} = G(\mathbf{I}_{y_o} | \mathbf{y}_f)$, that is, the generator has no effect. To prevent this situation, we regularize the mask with a l_2 -weight penalty. Also, to enforce smooth spatial color transformation when combining the pixel from the input image and the color transformation \mathbf{C} , we perform a *Total Variation Regularization* over \mathbf{A} . The attention loss $\mathcal{L}_A(G, \mathbf{I}_{y_o}, \mathbf{y}_f)$ can therefore be defined as:

$$\lambda_{\text{TV}} \mathbb{E}_{\mathbf{I}_{y_o} \sim \mathbb{P}_o} \left[\sum_{i,j}^{H,W} [(\mathbf{A}_{i+1,j} - \mathbf{A}_{i,j})^2 + (\mathbf{A}_{i,j+1} - \mathbf{A}_{i,j})^2] \right] + \mathbb{E}_{\mathbf{I}_{y_o} \sim \mathbb{P}_o} [\|\mathbf{A}\|_2] \quad (2)$$

where $\mathbf{A} = G_A(\mathbf{I}_{y_o} | \mathbf{y}_f)$ and $\mathbf{A}_{i,j}$ is the i, j entry of \mathbf{A} . λ_{TV} is a penalty coefficient.

Conditional Expression Loss. While reducing the *image adversarial loss*, the generator must also reduce the error produced by the AUs regression head on top of D . In this way, G not only learns to render realistic samples but also learns to satisfy the target facial expression encoded by \mathbf{y}_f . This loss is defined with two components: an AUs regression loss with fake images used to optimize G , and an AUs regression loss of real images used to learn the regression head on top of D . This loss $\mathcal{L}_Y(G, D_Y, \mathbf{I}_{y_o}, \mathbf{y}_o, \mathbf{y}_f)$ is computed as:

$$\mathbb{E}_{\mathbf{I}_{y_o} \sim \mathbb{P}_o} [\|D_Y(G(\mathbf{I}_{y_o} | \mathbf{y}_f)) - \mathbf{y}_f\|_2^2] + \mathbb{E}_{\mathbf{I}_{y_o} \sim \mathbb{P}_o} [\|D_Y(\mathbf{I}_{y_o}) - \mathbf{y}_o\|_2^2]. \quad (3)$$

Identity Loss. With the previously defined losses the generator is enforced to generate photo-realistic face transformations. However, without ground-truth supervision, there is no constraint to guarantee that the face in both the input and output images correspond to the same person. Using a *cycle consistency*

loss [38] we force the generator to maintain the identity of each individual by penalizing the difference between the original image \mathbf{I}_{y_o} and its reconstruction:

$$\mathcal{L}_{\text{idt}}(G, \mathbf{I}_{y_o}, \mathbf{y}_o, \mathbf{y}_f) = \mathbb{E}_{\mathbf{I}_{y_o} \sim \mathbb{P}_o} [\|G(G(\mathbf{I}_{y_o}|\mathbf{y}_f)|\mathbf{y}_o) - \mathbf{I}_{y_o}\|_1]. \quad (4)$$

To produce realistic images it is critical for the generator to model both low and high frequencies. Our *PatchGan* based critic D_I already enforces high-frequency correctness by restricting our attention to the structure in local image patches. To also capture low-frequencies it is sufficient to use l_1 -norm. In preliminary experiments, we also tried replacing l_1 -norm with a more sophisticated *Perceptual* [11] loss, although we did not observe improved performance.

Full Loss. To generate the target image \mathbf{I}_{y_g} , we build a loss function \mathcal{L} by linearly combining all previous partial losses:

$$\begin{aligned} \mathcal{L} = & \mathcal{L}_I(G, D_I, \mathbf{I}_{y_r}, \mathbf{y}_g) + \lambda_y \mathcal{L}_Y(G, D_Y, \mathbf{I}_{y_r}, \mathbf{y}_r, \mathbf{y}_g) \\ & + \lambda_A (\mathcal{L}_A(G, \mathbf{I}_{y_g}, \mathbf{y}_r) + \mathcal{L}_A(G, \mathbf{I}_{y_r}, \mathbf{y}_g)) + \lambda_{\text{idt}} \mathcal{L}_{\text{idt}}(G, \mathbf{I}_{y_r}, \mathbf{y}_r, \mathbf{y}_g), \end{aligned} \quad (5)$$

where λ_A , λ_y and λ_{idt} are the hyper-parameters that control the relative importance of every loss term. Finally, we can define the following minimax problem:

$$G^* = \arg \min_G \max_{D \in \mathcal{D}} \mathcal{L}, \quad (6)$$

where G^* draws samples from the data distribution. Additionally, we constrain our discriminator D to lie in \mathcal{D} , that represents the set of 1-Lipschitz functions.

5 Implementation Details

Our generator builds upon the variation of the network from Johnson *et al.* [11] proposed by [38] as it proved to achieve impressive results for image-to-image mapping. We have slightly modified it by substituting the last convolutional layer with two parallel convolutional layers, one to regress the color mask \mathbf{C} and the other to define the attention mask \mathbf{A} . We also observed that changing batch normalization in the generator by instance normalization improved training stability. For the critic we have adopted the *PatchGan* architecture of [10], but removing feature normalization. Otherwise, when computing the gradient penalty, the norm of the critic’s gradient would be computed with respect to the entire batch and not with respect to each input independently.

The model is trained on the EmotioNet dataset [3]. We use a subset of 200,000 samples (over 1 million) to reduce training time. We use Adam [14] with learning rate of 0.0001, beta1 0.5, beta2 0.999 and batch size 25. We train for 30 epochs and linearly decay the rate to zero over the last 10 epochs. Every 5 optimization steps of the critic network we perform a single optimization step of the generator. The weight coefficients for the loss terms in Eq. (5) are set to $\lambda_{\text{gp}} = 10$, $\lambda_A = 0.1$, $\lambda_{\text{TV}} = 0.0001$, $\lambda_y = 4000$, $\lambda_{\text{idt}} = 10$. To improve stability we tried updating the critic using a buffer with generated images in different updates of the generator as proposed in [32] but we did not observe performance improvement. The model takes two days to train with a single GeForce[®] GTX 1080 Ti GPU.




Fig. 4. Single AUs edition. Specific AUs are activated at increasing levels of intensity (from 0.33 to 1). The first row corresponds to a zero intensity application of the AU which correctly produces the original image in all cases.

6 Experimental Evaluation

This section provides a thorough evaluation of our system. We first test the main component, namely the single and multiple AUs editing. We then compare our model against current competing techniques in the task of discrete emotions editing and demonstrate our model’s ability to deal with images in the wild and its capability to generate a wide range of anatomically coherent face transformations. Finally, we discuss the model’s limitations and failure cases.

It is worth noting that in some of the experiments the input faces are not cropped. In this cases we first use a detector ² to localize and crop the face, apply the expression transformation to that area with Eq. (1), and finally place the generated face back to its original position in the image. The attention mechanism guaranties a smooth transition between the morphed cropped face and the original image. As we shall see later, this three steps process results on higher resolution images compared to previous models. Supplementary material can be found on <http://www.albertpumarola.com/research/GANimation/>.

6.1 Single Action Units Edition

We first evaluate our model’s ability to activate AUs at different intensities while preserving the person’s identity. Figure 4 shows a subset of 9 AUs individually transformed with four levels of intensity (0, 0.33, 0.66, 1). For the case of 0 intensity it is desired not to change the corresponding AU. The model properly handles this situation and generates an identical copy of the input image for every case. The ability to apply an identity transformation is essential to ensure that non-desired facial movement will not be introduced.

² We use the face detector from https://github.com/ageitgey/face_recognition.



Fig. 5. Attention Model. Details of the intermediate attention mask **A** (first row) and the color mask **C** (second row). The bottom row images are the synthesized expressions. Darker regions of the attention mask **A** show those areas of the image more relevant for each specific AU. Brighter areas are retained from the original image.

For the non-zero cases, it can be observed how each AU is progressively accentuated. Note the difference between generated images at intensity 0 and 1. The model convincingly renders complex facial movements which in most cases are difficult to distinguish from real images. It is also worth mentioning that the independence of facial muscle cluster is properly learned by the generator. AUs relative to the eyes and half-upper part of the face (AUs 1, 2, 4, 5, 45) do not affect the muscles of the mouth. Equivalently, mouth related transformations (AUs 10, 12, 15, 25) do not affect eyes nor eyebrow muscles.

Fig. 5 displays, for the same experiment, the attention **A** and color **C** masks that produced the final result \mathbf{I}_{y_g} . Note how the model has learned to focus its attention (darker area) onto the corresponding AU in an unsupervised manner. In this way, it relieves the color mask from having to accurately regress each pixel value. Only the pixels relevant to the expression change are carefully estimated, the rest are just noise. For example, the attention is clearly obviating background pixels allowing to directly copy them from the original image. This is a key ingredient to later being able to handle images in the wild (see Section 6.5).

6.2 Simultaneous Edition of Multiple AUs

We next push the limits of our model and evaluate it in editing multiple AUs. Additionally, we also assess its ability to interpolate between two expressions. The results of this experiment are shown in Fig. 1, the first column is the original image with expression \mathbf{y}_r , and the right-most column is a synthetically generated image conditioned on a target expression \mathbf{y}_g . The rest of columns result from evaluating the generator conditioned with a linear interpolation of the original and target expressions: $\alpha\mathbf{y}_g + (1 - \alpha)\mathbf{y}_r$. The outcomes show a very remarkable smooth and consistent transformation across frames. We have intentionally selected challenging samples to show robustness to light conditions and even, as in the case of the avatar, to non-real world data distributions which were not previously seen by the model. These results are encouraging to further extend the model to video generation in future works.




Fig. 6. Qualitative comparison with state-of-the-art. Facial Expression Synthesis results for: DIAT [20], CycleGAN [28], IcGAN [26] and StarGAN [4]; and ours. In all cases, we represent the input image and seven different facial expressions. As it can be seen, our solution produces the best trade-off between visual accuracy and spatial resolution. Some of the results of StarGAN, the best current approach, show certain level of blur. Images of previous models were taken from [4].

6.3 Discrete Emotions Editing

We next compare our approach, against the baselines DIAT [20], CycleGAN [28], IcGAN [26] and StarGAN [4]. For a fair comparison, we adopt the results of these methods trained by the most recent work, StarGAN, on the task of rendering discrete emotions categories (e.g., happy, sad and fearful) in the RaFD dataset [16]. Since DIAT [20] and CycleGAN [28] do not allow conditioning, they were independently trained for every possible pair of source/target emotions. We next briefly discuss the main aspects of each approach:

DIAT [20]. Given an input image $x \in X$ and a reference image $y \in Y$, DIAT learns a GAN model to render the attributes of domain Y in the image x while conserving the person’s identity. It is trained with the classic *adversarial loss* and a *cycle loss* $\|x - G_{Y \rightarrow X}(G_{X \rightarrow Y}(x))\|_1$ to preserve the person’s identity.

CycleGAN [28]. Similar to DIAT [20], CycleGAN also learns the mapping between two domains $X \rightarrow Y$ and $Y \rightarrow X$. To train the domain transfer, it uses a regularization term denoted *cycle consistency loss* combining two cycles: $\|x - G_{Y \rightarrow X}(G_{X \rightarrow Y}(x))\|_1$ and $\|y - G_{X \rightarrow Y}(G_{Y \rightarrow X}(y))\|_1$.

IcGAN [26]. Given an input image, IcGAN uses a pretrained encoder-decoder to encode the image into a latent representation in concatenation with an expression vector \mathbf{y} to then reconstruct the original image. It can modify the expression by replacing \mathbf{y} with the desired expression before going through the decoder.




Fig. 7. Sampling the face expression distribution space. As a result of applying our AU-parametrization through the vector \mathbf{y}_g , we can synthesize, from the same source image \mathbf{I}_{v_r} , a large variety of photo-realistic images.

StarGAN [4]. An extension of *cycle loss* for simultaneously training between multiple datasets with different data domains. It uses a mask vector to ignore unspecified labels and optimize only on known ground-truth labels. It yields more realistic results when training simultaneously with multiple datasets.

Our model differs from these approaches in two main aspects. First, we do not condition the model on discrete emotions categories, but we learn a basis of anatomically feasible warps that allows generating a continuum of expressions. Secondly, the use of the attention mask allows applying the transformation only on the cropped face, and put it back onto the original image without producing any artifact. As shown in Fig. 6, besides estimating more visually compelling images than other approaches, this results on images of higher spatial resolution.

6.4 High Expressions Variability

Given a single image, we next use our model to produce a wide range of anatomically feasible face expressions while conserving the person’s identity. In Fig. 7 all faces are the result of conditioning the input image in the top-left corner with a desired face configuration defined by only 14 AUs. Note the large variability of anatomically feasible expressions that can be synthesized with only 14 AUs.

6.5 Images in the Wild

As previously seen in Fig. 5, the attention mechanism not only learns to focus on specific areas of the face but also allows merging the original and generated image background. This allows our approach to be easily applied to images in the wild while still obtaining high resolution images. For these images we follow the



Fig. 8. Qualitative evaluation on images in the wild. **Top:** We represent an image (left) from the film “*Pirates of the Caribbean*” and an its generated image obtained by our approach (right). **Bottom:** In a similar manner, we use an image frame (left) from the series “*Game of Thrones*” to synthesize five new images with different expressions.

detection and cropping scheme we described before. Fig. 8 shows two examples on these challenging images. Note how the attention mask allows for a smooth and unnoticeable merge between the entire frame and the generated faces.

6.6 Pushing the Limits of the Model

We next push the limits of our network and discuss the model limitations. We have split success cases into six categories which we summarize in Fig. 9-top. The first two examples (top-row) correspond to human-like sculptures and non-realistic drawings. In both cases, the generator is able to maintain the artistic effects of the original image. Also, note how the attention mask ignores artifacts such as the pixels occluded by the glasses. The third example shows robustness to non-homogeneous textures across the face. Observe that the model is not trying to homogenize the texture by adding/removing the beard’s hair. The middle-right category relates to anthropomorphic faces with non-real textures. As for the Avatar image, the network is able to warp the face without affecting its texture. The next category is related to non-standard illuminations/colors for which the model has already been shown robust in Fig. 1. The last and most surprising category is face-sketches (bottom-right). Although the generated face suffers from some artifacts, it is still impressive how the proposed method is still capable of finding sufficient features on the face to transform its expression from worried to excited. The second case shows failures with non-previously seen occlusions such as an eye patch causing artifacts in the missing face attributes.

We have also categorized the failure cases in Fig. 9-bottom, all of them presumably due to insufficient training data. The first case is related to errors in the attention mechanism when given extreme input expressions. The attention does not weight sufficiently the color transformation causing transparencies.

The model also fails when dealing with non-human anthropomorphic distributions as in the case of cyclopes. Lastly, we tested the model behavior when dealing with animals and observed artifacts like human face features.

References

1. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein GAN. arXiv preprint arXiv:1701.07875 (2017)
2. Baltrušaitis, T., Mahmoud, M., Robinson, P.: Cross-dataset learning and person-specific normalisation for automatic action unit detection. In: FG (2015)
3. Benitez-Quiroz, C.F., Srinivasan, R., Martinez, A.M., et al.: Emotionet: An accurate, real-time algorithm for the automatic annotation of a million facial expressions in the wild. In: CVPR (2016)
4. Choi, Y., Choi, M., Kim, M., Ha, J.W., Kim, S., Choo, J.: Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. CVPR (2018)
5. Du, S., Tao, Y., Martinez, A.M.: Compound facial expressions of emotion. Proceedings of the National Academy of Sciences p. 201322355 (2014)
6. Ekman, P., Friesen, W.: Facial action coding system: A technique for the measurement of facial movement. Consulting Psychologists Press (1978)
7. Fischler, M.A., Elschlager, R.A.: The representation and matching of pictorial structures. IEEE Transactions on Computers **22**(1), 67–92 (1973)
8. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NIPS (2014)
9. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of wasserstein GANs. In: NIPS (2017)
10. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR (2017)
11. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: ECCV (2016)
12. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of GANs for improved quality, stability, and variation. In: ICLR (2018)
13. Kim, T., Cha, M., Kim, H., Lee, J., Kim, J.: Learning to discover cross-domain relations with generative adversarial networks. In: ICML (2017)
14. Kingma, D., Ba, J.: ADAM: A method for stochastic optimization. In: ICLR (2015)
15. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: ICLR (2014)
16. Langner, O., Dotsch, R., Bijlstra, G., Wigboldus, D.H., Hawk, S.T., Van Knippenberg, A.: Presentation and validation of the radboud faces database. Cognition and emotion **24**(8), 1377–1388 (2010)
17. Larsen, A.B.L., Sønderby, S.K., Larochelle, H., Winther, O.: Autoencoding beyond pixels using a learned similarity metric. In: ICML (2016)
18. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: CVPR (2017)
19. Li, C., Wand, M.: Precomputed real-time texture synthesis with markovian generative adversarial networks. In: ECCV (2016)
20. Li, M., Zuo, W., Zhang, D.: Deep identity-aware transfer of facial attributes. arXiv preprint arXiv:1610.05586 (2016)
21. Liu, M.Y., Breuel, T., Kautz, J.: Unsupervised image-to-image translation networks. In: NIPS (2017)
22. Mathieu, M., Couprie, C., LeCun, Y.: Deep multi-scale video prediction beyond mean square error. In: ICLR (2016)
23. Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784 (2014)

24. Odena, A., Olah, C., Shlens, J.: Conditional image synthesis with auxiliary classifier GANs. In: ICML (2017)
25. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: CVPR (2016)
26. Perarnau, G., van de Weijjer, J., Raducanu, B., Álvarez, J.M.: Invertible conditional GANs for image editing. arXiv preprint arXiv:1611.06355 (2016)
27. Pumarola, A., Agudo, A., Sanfeliu, A., Moreno-Noguer, F.: Unsupervised person image synthesis in arbitrary poses. In: CVPR (2018)
28. Radford, A., Metz, L., Chintala, S.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICLR (2016)
29. Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., Lee, H.: Generative adversarial text to image synthesis. In: ICML (2016)
30. Scherer, K.R.: Emotion as a process: Function, origin and regulation. *Social Science Information* **21**, 555–570 (1982)
31. Shen, W., Liu, R.: Learning residual images for face attribute manipulation. In: CVPR (2017)
32. Shrivastava, A., Pfister, T., Tuzel, O., Susskind, J., Wang, W., Webb, R.: Learning from simulated and unsupervised images through adversarial training. In: CVPR (2017)
33. Wang, X., Gupta, A.: Generative image modeling using style and structure adversarial networks. In: ECCV (2016)
34. Wang, Z., Liu, D., Yang, J., Han, W., Huang, T.: Deep networks for image super-resolution with sparse prior. In: ICCV (2015)
35. Yu, H., Garrod, O.G., Schyns, P.G.: Perception-driven facial expression synthesis. *Computers & Graphics* **36**(3) (2012)
36. Zafeiriou, S., Trigeorgis, G., Chrysos, G., Deng, J., Shen, J.: The menpo facial landmark localisation challenge: A step towards the solution. In: CVPRW (2017)
37. Zhang, H., Xu, T., Li, H., Zhang, S., Huang, X., Wang, X., Metaxas, D.: Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: ICCV (2017)
38. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV (2017)
39. Zhu, S., Fidler, S., Urtasun, R., Lin, D., Loy, C.C.: Be your own prada: Fashion synthesis with structural coherence. In: ICCV (2017)