

# 3D Dense Geometry-Guided Facial Expression Synthesis by Adversarial Learning

Rumeysa Bodur<sup>1</sup>, Binod Bhattarai<sup>1</sup>, and Tae-Kyun Kim<sup>1,2</sup>

<sup>1</sup>Imperial College London, UK

<sup>2</sup>KAIST, South Korea

{r.bodur18, b.bhattarai, tk.kim}@imperial.ac.uk

## Abstract

Manipulating facial expressions is a challenging task due to fine-grained shape changes produced by facial muscles and the lack of input-output pairs for supervised learning. Unlike previous methods using Generative Adversarial Networks (GAN), which rely on cycle-consistency loss or sparse geometry (landmarks) loss for expression synthesis, we propose a novel GAN framework to exploit 3D dense (depth and surface normals) information for expression manipulation. However, a large-scale dataset containing RGB images with expression annotations and their corresponding depth maps is not available. To this end, we propose to use an off-the-shelf state-of-the-art 3D reconstruction model to estimate the depth and create a large-scale RGB-Depth dataset after a manual data clean-up process. We utilise this dataset to minimise the novel depth consistency loss via adversarial learning (note we do not have ground truth depth maps for generated face images) and the depth categorical loss of synthetic data on the discriminator. In addition, to improve the generalisation and lower the bias of the depth parameters, we propose to use a novel confidence regulariser on the discriminator side of the framework. We extensively performed both quantitative and qualitative evaluations on two publicly available challenging facial expression benchmarks: AffectNet and RaFD. Our experiments demonstrate that the proposed method outperforms the competitive baseline and existing arts by a large margin.

## 1. Introduction

Face expression manipulation is an active and challenging research problem [3, 24, 4]. It has several applications such as movie industry and e-commerce. Unlike the manipulation of other facial attributes [3, 16, 1], e.g. *hair/skin*




Figure 1. **Sparse Geometry vs Dense Geometry.** Some of the randomly selected images from AffectNet and their corresponding geometric information. We can observe that sparse geometry, i.e. landmarks, is incapable of capturing expression specific fine details, e.g. wrinkles on the cheek and jaw area, nose and eyebrow shapes.

*colour, wearing glasses, beard*, expression manipulation is relatively fine-grained in nature. The change in expressions is caused by geometrical actions of facial muscles in a detailed manner. Figure 1 shows images with two different expressions, sad and happy. From these images we can observe that different expressions bear different geometric information in a fine-grained manner. The task of expression manipulation becomes even more challenging when source and target expression pairs are not available to train the model, which is the case in this study.

In one of the earliest studies [5], the Facial Action Coding System (FACS) was developed for describing facial expressions, which is popularly known as Action Units (AUs). These AUs are anatomically related to the contractions and relaxations of certain facial muscles. Although the number of AUs is only 30, more than 7000 different combinations of AUs that describe expressions have been observed [27]. This makes the task more challenging than other coarse attribute/image manipulation tasks, e.g. higher order face attribute editing [3], season translation [16]. These kind of

problems cannot be addressed only by optimising existing loss functions, such as cycle consistency [3] or by adding a target attribute classifier on the discriminator side [3, 8].

A recent study on GAN [24] proposes to use AUs to manipulate expressions and shows a performance improvement over prior-arts, which include CycleGAN [35], DIAT [15] and IcGAN [22]. These methods do not explicitly address such details in their objective functions, but rather only deal with coarse characteristics, and this ultimately results in a sub-optimal solution. The main drawback of notable existing methods exploiting geometry, such as GANimation [24] (action units) and Gannotation [24] (sparse-landmarks), is the use of annotations which are sparse in nature. As presented in Figure 1, these sparse geometrical annotations fail to capture the geometries of different expressions precisely. To address these issues, we propose to introduce a dense 3D depth consistency loss for expression manipulation. The use of dense 3D geometry guides how different expressions should appear in image pixel space and thereby the process of image synthesis by translating from one expression to another.

RGB-Depth consistency loss has been successfully applied for the semantic segmentation task [2], which is formulated as a domain translation problem, i.e. from simulation to real world. However, this task is less demanding than ours since the label as well as the geometric information of the translated image remain the same as that of the source. On the other hand, our objective is to manipulate an image from the source expression to a target expression involving faithful transformation of expression relevant geometry without distorting the identity information. Also note that unlike previous works [4, 26, 13, 30, 25] on expression manipulation utilising depth information, our work is a target-label conditioned framework where the target expression is a simple one-hot encoded vector. Hence, it does not require source-target pairs during both training and testing time. Although, constraining depth information on model parameters are effective, a reasonably good size of RGB-D dataset for expressions is not available to date. Among the existing datasets, EURECOM Kinect face data set [17] consists of 3 expressions (neutral, smiling, open mouth) of 50 subjects, and consisting of 2500 images, BU-3DFE [34] provides a database with 8 expressions of 100 subjects. These datasets are too small to train a generative model optimally.

To overcome the challenge of the unavailability of RGB-D pairs in a scale to train a GAN, we propose to use an off-the-shelf 3D face reconstruction model [31] and extract the depth and surface normal information from the reconstructions. As discussed, since in our case the category of the image changes when translated from source to target domain, we propose to minimise the target attribute cross entropy loss on depth domain to ensure that the depth of the reconstructed image is also consistent with the target la-

bels. As we are aware, the depth images extracted from the off-the-shelf face 3D reconstruction model are error-prone and sub-optimal for our purpose. To address this challenge, we propose to penalise the depth estimator by employing a confidence penalty [23]. This regulariser is effective to improve the accuracy of classification models trained on noisy labels [10].

Our contributions can be summarised as follows:

- We propose a novel geometric consistency loss to guide the expression synthesis process using depth and surface normals information.
- We estimated the depth and surface normal parameters for two challenging expression benchmarks, namely AffectNet and RaFD, and will make these annotations public upon the acceptance of the paper.
- We propose a novel regularisation on the discriminator to penalise the confidence of the depth estimator in order to improve the generalisation of the cross-data set model parameters.
- We evaluated the proposed method on two challenging benchmarks. Our experiments demonstrate that the proposed method is superior to the competitive baselines both in quantitative and qualitative metrics.

## 2. Related Work

**Unpaired Image to Image Translation.** One of the applications of image-to-image translation by Generative Adversarial Networks (GANs) [35, 11] is the synthesis of attributes and facial expressions on given face images. Various methods have been proposed based on the GAN framework in order to accomplish this task. In general, the training sets used for image-to-image translation problems consist of aligned image pairs. However, since paired datasets for various facial expressions and attributes are quite small and constructed in a controlled environment, most of the methods are designed for datasets with unpaired images. CycleGAN [35] tackles the problem of unpaired image-to-image translation by coupling the mapping learned by the generator with an inverse mapping and an additional cycle consistency loss. This approach is employed by most of the related studies in order to preserve key attributes between the input and output images, and thus the person’s identity in the given source image. In StarGAN [3], the multi-domain image-to-image translation problem is approached by using a single generator instead of training a separate generator for each domain pair. In AttGAN [8] an attribute classification constraint is applied on the generated image rather than imposing an attribute-independence constraint on the latent encoding. The method is further extended to generate attributes of different styles by introducing a set of

style controllers and a style predictor. Unlike other attribute generation methods ELEGANT [33] aims to generate from exemplars, and thus uses latent encoding instead of labels. For multiple attribute generation, they encode the attributes in a disentangled manner, and for increasing the quality they adopt residual learning and multi-scale discriminators.

**Geometry Guided Expression Synthesis.** Existing works incorporating geometric information yields better quality of synthetic images. GP-GAN [4] generates faces from only sparse 2D landmarks. Similarly, GANnotation [26] is a facial-landmark guided model that simultaneously changes the pose and the expression. In addition, this method also minimises triple consistency loss for bridging the gap between the distributions of the input and generated image domains. G2GAN [30] is also another recent work on expression manipulation guided by 2D facial landmarks. In GAGAN [13], latent variables are sampled from a probability distribution of a statistical shape model i.e. mean and eigenvectors of landmark vectors, and the generated output is mapped into a canonical coordinate frame by using a differentiable piece-wise affine warping.

In GC-GAN [25], they apply contrastive learning in GAN for expression transfer. Their network consists of two encoder-decoder networks and a discriminator network. One encodes the image whilst the other encodes the landmarks of the target and the source disentangled from the image. GANimation [24] manages to generate a continuum of facial expression by conditioning the GAN on a 1-dimensional vector that indicates the presence and magnitude of AUs rather than just category labels. It benefits from both cycle loss and geometric loss, however the annotations are sparse. Whilst the studies mentioned above mostly rely on sparse geometric information, some recent studies go beyond landmarks to exploit geometric information and make use of surface normals. [28] exploits surface normals information in order to address the coarse attribute manipulation task whereas in our work we deal with the more challenging expression manipulation task. Similarly, [19] feeds the surface normal maps of the target as a condition to the networks for modifying the expression of the given image. Unlike this work, we are dealing with an unpaired image translation problem, where the expressions is given only as a categorical label.

### 3. Proposed Method

The schematic diagram of our method can be seen in Figure 2. As seen in the Figure, our pipeline consists of two generative adversarial networks in a cascaded form, namely the RGB adversarial network and the depth adversarial network. One of our contributions lies in introducing a depth network to the RGB network and training the framework in an end-to-end fashion. As our RGB adversarial network, we choose StarGAN [3], which is one of the representa-

tive conditional GAN architectures of contemporary frameworks that has been widely used for unpaired multi-domain image translation. Since our method is generic, it can be employed in any other RGB adversarial network. In this section, we will explain our RGB-Depth dataset construction followed by architectural details and the training process.

#### 3.1. RGB-Depth Dataset Creation

To train a generator for estimating depth from RGB images accurately, there is a need for a large-scale dataset consisting of RGB image and depth map pairs of faces with different expressions. As discussed in Section 1, there is no such large-scale dataset available that fits the requirement to train a GAN adequately. To fill this gap, we propose to create a large-scale RGB-depth dataset. Hence, we utilise the publicly available model of [31] to reconstruct 3D meshes of a subset of images from existing datasets, AffectNet[18] and RaFD[14], and augment them with depth maps. However, some of the 3D meshes were poor in quality, which we carefully discarded manually. We projected the reconstructed meshes and interpolated missing parts of the projections in order to acquire the corresponding depth maps and normal maps. Samples obtained by this process are shown in Figure 3.

To validate the generalisation and statistical significance of the RGB image and depth map pairs, we use the datasets we constructed to train a model in an adversarial manner, which generates depth maps and normal maps from RGB images. As estimating depth from a single RGB image is an ill-posed problem, it is essential to take the global scene into account. Since we need local as well as global consistency on the predicted depth maps, we train the model in an adversarial manner [7]. Although recent CNNs are capable of understanding the global relationship too, their objective functions minimise the combination of per-pixel wise losses. To differentiate the predicted depth maps from the extracted depth maps, for convenience, we refer to the extracted depth maps as *ground truth depth maps*. Figure 3 shows the predicted and ground truth depth maps and normal maps of some of the examples from the real test set. From these examples, we observe that predicted depth maps and normal maps are quite similar to the ground truths. This demonstrates that the model trained with our dataset generalises well to unseen examples. Hence, it shows the potential uses of our dataset for future use cases as well. We will make our dataset *publicly available* for the research community upon the acceptance of this paper. We name the depth augmented AffectNet and RaFD datasets as AffectNet-D and RaFD-D, respectively, for future use.

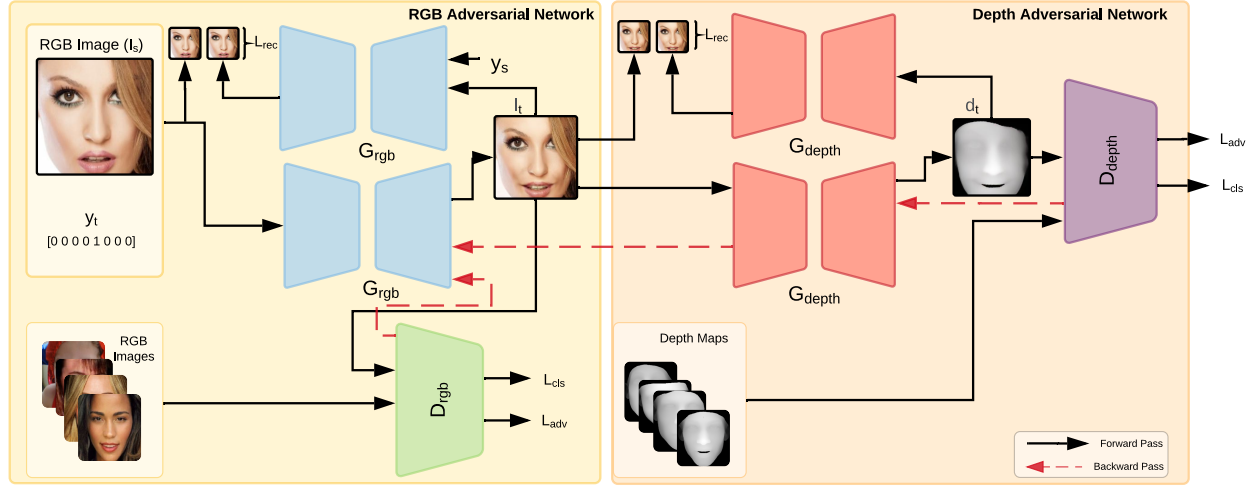


Figure 2. **Overview of the Proposed End-to-End Pipeline.** It consists of two generative adversarial networks, namely the RGB network and the depth network.  $G_{rgb}$  takes as input the RGB source image along with the target label and generates a synthetic image with the with the given target expression. This image is passed to  $G_{rgb}$  to generate its estimated depth and thereby to calculate the depth consistency loss. This loss helps capturing detailed 3D geometric information, which is then aligned in generated RGB images.

### 3.2. Model

We have a scenario in which an RGB input image and a target expression label in the form of a one-hot encoded vector are given. Our objective is to translate this image accurately to the given target label with high-quality measurements. To this end, we propose to train two types of adversarial networks, namely RGB and Depth, in an end-to-end fashion. In the following part, we discuss about these networks in detail.

**RGB Adversarial Network.** As seen from Figure 2, the first part of the pipeline describes the RGB adversarial network that consists of a generator  $G_{rgb}$  and a discriminator  $D_{rgb}$ . This part is StarGAN [3], however we would like to emphasise that our method is not constrained to one architecture. The discriminator serves two functions  $D_{rgb} : \{D_{adv}^{rgb}, D_{cls}^{rgb}\}$ , with  $D_{adv}^{rgb}$  providing a probability distribution over real and fake domain, while  $D_{cls}^{rgb}$  provides a distribution over the expression classes. The network takes as input an RGB source image,  $I_s$ , along with a target expression label,  $y_t$ , as condition.  $G_{rgb}$  translates the source image in the guidance of the given target label,  $y_t$ , yielding a synthetic image  $I_t$ . Then,  $D_{adv}^{rgb}$ , which is a discriminator trained to distinguish between real and synthetic RGB images, is applied on the generated image to ensure a realistic look in the target domain. The adversarial loss used for training  $D_{rgb}$  is given in Eqn. 1.

$$L_{adv}^{rgb} = \mathbb{E}_{I_s} [\log(D_{adv}^{rgb}(I_s))] + \mathbb{E}_{I_s, y_t} [\log(1 - D_{adv}^{rgb}(G_{rgb}(I_s, y_t)))] \quad (1)$$

The discriminator contains an auxiliary classifier,  $D_{cls}^{rgb}$ , which guides  $G_{rgb}$  to generate images that can be confi-

dently classified as the target expression. We minimise a categorical cross-entropy loss as follows:

$$L_{cls}^{rgb} = \mathbb{E}_{I_s, y_t} [\log(D_{cls}^{rgb}(y_t | G_{rgb}(I_s, y_t)))] \quad (2)$$

To preserve attribute excluding details, a reconstruction loss is employed, which is a cycle consistency loss and is formulated in Eqn.3. Here,  $y_s$  represents the original label of the source image  $I_s$ .

$$L_{rec}^{rgb} = \mathbb{E}_{(I_s, y_s), y_t} [\|I_s - G_{rgb}(G_{rgb}(I_s, y_t), y_s)\|_1] \quad (3)$$

The overall objective to be minimised for the RGB adversarial network can be summed as follows:

$$L_{rgb} = \lambda_{adv} L_{adv}^{rgb} + \lambda_{cls} L_{cls}^{rgb} + \lambda_{rec} L_{rec}^{rgb} \quad (4)$$

In Equation 4,  $\lambda_{adv}$ ,  $\lambda_{cls}$ , and  $\lambda_{rec}$  are hyper-parameters to control the weight of adversarial, classification, and reconstruction losses, and are set by cross-validation. From Equation 4, we can see that none of these objectives are explicitly modeling the fine-grained details of the expressions. Thus, the solution from RGB adversarial network alone remains sub-optimal.

**Depth Adversarial Network.** To address the limitations of the RGB adversarial network, we propose to append another network that operates on the depth domain and guides the RGB network, because depth maps are able to capture the geometric information of various expressions. As depicted in the right-half of Figure 2, similar to the RGB adversarial network, the depth network also consists of a generator  $G_{depth}$  and a discriminator  $D_{depth} : \{D_{adv}^{depth}, D_{cls}^{depth}\}$ . In this case the generator,  $G_{depth}$ , takes as input the synthetic image,  $I_t$ , generated by  $G_{rgb}$  and generates the estimated




Figure 3. **RGB Images and Their Ground Truth/Predicted Depth and Surface Normal Maps.** Real test images with their corresponding ground truth depth maps extracted from the off-the-shelf 3D model and the depth maps predicted by our depth-estimator network. We can clearly see that the depth maps are quite aligned with the geometry of real images, and our network trained on the training set of image-depth pairs is able to preserve the geometry.

depth map,  $d_t$ , which is then passed to the discriminator.  $D_{depth}$  is trained to distinguish between realistic and synthetic depth maps by using the following adversarial loss function, where  $d_s$  represents the real depthmaps:

$$L_{adv}^{depth} = \mathbb{E}_{d_s} [\log(D_{adv}^{depth}(d_s))] + \mathbb{E}_{I_t} [\log(1 - D_{adv}^{depth}(G_{depth}(I_t)))] \quad (5)$$

Again, an auxiliary classifier,  $D_{cls}^{depth}$ , is added on top of the discriminator in order to generate depth maps that can be confidently classified as the target expression,  $y_t$ . The loss function used for this classifier is given in Eqn. 6.

$$L_{cls}^{depth} = \mathbb{E}_{I_t, y_t} [\log(D_{cls}^{depth}(y_t | G_{depth}(I_t)))] \quad (6)$$

The Depth Network is trained in a similar way to the multi-domain translation network in [3], where the domain is given as a condition. Thus, the network is able to translate from RGB to depth domain as well as from depth to RGB domain. We denote the latter translation as  $G'_{depth}$ . This network takes the estimated depth map and generates an RGB image. As done in the RGB adversarial network, we employ the reconstruction loss given in Eqn. 7.

$$L_{rec}^{depth} = \mathbb{E}_{I_t} [\|I_t - G'_{depth}(G_{depth}(I_t))\|_1] \quad (7)$$

The losses used for the depth network can be summed as follows:

$$L_{depth} = \lambda_{adv} L_{adv}^{depth} + \lambda_{cls} L_{cls}^{depth} + \lambda_{rec} L_{rec}^{depth} \quad (8)$$

where  $\lambda_{adv}$ ,  $\lambda_{cls}$ , and  $\lambda_{rec}$  are hyper-parameters for adversarial, classification and reconstruction loss. We set their values by cross-validation. Minimising dense depth loss which encodes fine-grained geometric information related to expressions helps to synthesise high-quality expression manipulated images.

**Overall Training Objective.** The overall objective function used for the end-to-end training of the networks is as follows:

$$\min_{\{G_{rgb}, G_{depth}\}} \max_{\{D_{rgb}, D_{depth}\}} (L_{rgb} + L_{depth})$$

From Eqns. 5, 6 and 7 we can see that the loss incurred by the synthetic image on the depth domain is propagated to the RGB adversarial network. This will ultimately guide the RGB generator to synthesise the image with the target attributes in such a way that their geometric information is also consistent.

**Pre-training the Depth Network.** The Depth network that we employed in this framework is trained offline on the data set augmented with the depth maps. Please refer to Section 3.1 for creation and augmentation of the depth maps on AffectNet and RaFD. We train the Depth Network with strong supervision in prior utilising the RAFD-D and AffectNet-D datasets we constructed, so that it can generate reasonable depth maps for given synthetic inputs during the end-to-end training. Employing such pre-trained network to constrain the GAN training has been useful in other downstream tasks such as identity preservation [6]. Apart from the depth network losses described in the previous section, since we have paired data we employ a pixel loss and a perceptual loss. The pixel loss, as given in Eqn. 9, calculates the L1 distance between the depth map generated by  $G_{depth}$  for the input image,  $I_s$ , and its ground truth depth map,  $d_s$ :

$$L_{pix} = \mathbb{E}_{(I_s, d_s)} [\|d_s - G_{depth}(I_s)\|_1] \quad (9)$$

Instead of relying solely on the L1 or L2 distance, with the perceptual loss [12] defined in Eqn. 10 the model learns by using the error between high-level feature representations extracted by a pre-trained CNN, which in our case is VGG19 [29].

$$L_{per} = \mathbb{E}_{(I_s, d_s)} [\|V(d_s) - V(G_{depth}(I_s))\|_1] \quad (10)$$

Finally, to improve the generalisation of the depth network, we propose to introduce a regulariser on the depth

network called confidence penalty as shown in the Equation 11. This regulariser relaxes the confidence of the depth prediction model similar to regularisers that learn from noisy labels [10]. Such regularisers are successfully applied on image classification tasks [23]. As mentioned before, the depth maps are extracted from the off-the-self 3D reconstruction model and manual eradication of low quality depth maps has already made our data set ready to learn the model. Even then, this regulariser helps to further remove the noise from the data set and learn the robust model. In Eqn. 11  $H$  stands for the entropy and  $\beta$  controls the strength of this penalty. We set the parameters of  $\beta$  by cross-validation.

$$L_{cls}^{depth} = \mathbb{E}_{I_s, y_t} [\log(D_{cls}^{depth}(y_t | G_{depth}(I_s)))] - \beta H(D_{cls}^{depth}(y_t | G_{depth}(I_s))) \quad (11)$$

## 4. Experimental Results

**Datasets.** We extended two popular benchmarks, AffectNet [18] and RaFD [14], to AffectNet-D and RaFD-D. For AffectNet, we took the down sampled version, where a maximum of 15K images are selected for each class. This resulted in a reconstruction of 87K meshes for 8 expressions. RaFD[14] contains images of 67 models taken with 3 different gazes and 5 different camera angles, where each model shows 8 expressions. The images are aligned based on their landmarks and cropped to the size of  $256 \times 256$ .

**Implementation.** Our method is implemented on PyTorch [21]. We perform our experiments on a workstation with Intel i5-8500 3.0G, 32G memory, NVIDIA GTX1060 and NVIDIA GTX1080. As for the training schedule, we first pre-train the depth network for 200k epochs, then both GANs in an end-to-end fashion for 300k epochs, where the batch size is set to 15 and 8, respectively. We start with a learning rate of 0.0001 and decay it in every 1000 epochs. We update the discriminator 5 times per generator update. The models are trained with Adam optimiser.

**Evaluation metrics.** We evaluate our method both quantitatively and qualitatively. To quantitatively evaluate the generated images we compute the Frechet Inception Distance (FID)[9]. FID is a commonly used metric for assessing the quality and diversity of synthetic images. We further evaluate our method by applying a pre-trained face recognition model to calculate the identity loss for synthesised images. We also calculate the peak signal-to-noise ratio (PSNR) and the structural similarity (SSIM) [32] between the original and reconstructed images. Also, similar to the attribute generation rate reported on STGAN [16], we report expression generation rate in our experiments. To calculate it, we train a classifier independent of all models on the real training set and evaluate on the synthetic test sets.

**Compared Methods.** We compare our method to some of the state-of-the-art attribute manipulation baselines, Star-

GAN [3], IcGAN [22] and CycleGAN [35]. Among these methods we took StarGAN[3] as our baseline. We further compare our method to some recent studies that utilise geometric information, namely Ganimation [24] and Gannotation [26] (Please refer to Table 1). For Ganimation, we use their publicly available implementation and trained it on our dataset.

### 4.1. Quantitative Evaluation

We compare our method with existing arts on various quantitative metrics on both Affectnet and RaFD.

**Expression generation rate:** On AffectNet, we report the performance on 6 expressions excluding contempt and disgust, since these two classes are highly under-represented on this dataset. Apart from this accuracy score on AffectNet, all other quantitative results are based on all 8 expression classes for both datasets. As seen from Table 1, on AffectNet, when we apply our method on StarGAN, the performance increases by 3.8%. With the introduction of the confidence penalty to the depth network, this rate further increases by 1.3% yielding a 5.1% improvement over StarGAN by reaching 87.2%. Similarly, as seen from Table 2, for RaFD, with our method we get an accuracy of 62%, which corresponds to an 8.9% improvement over StarGAN.

**FID:** To assess the quality of the synthesised images we calculate the FID for both datasets. As seen in the tables, when our method is applied over StarGAN, the FID values decrease from 16.33 to 13.35 for AffectNet, and from 39.45 to 35.14 for RaFD. Since lower FID values indicate better image quality and diversity, these results verify that our method improves the synthetic images in both aspects.

**PSNR/SSIM:** We further compute the PSNR/SSIM scores to evaluate the reconstruction performance of the networks. Our method improves the PSNR score by 2.33 and 1.32 yielding a score of 35.29 and 32.57 for AffectNet and RaFD, respectively. Similarly, the SSIM score improves by 5% for AffectNet and 4% for RaFD resulting in 84% and 79%.

**Identity Preservation:** Finally, to assess how well the identity is preserved we employ the pre-trained model of VGG Face[20] to extract the features of real and synthesised images. We calculate the L2 loss between the features of each synthesised image and the corresponding real image. Figure 4 illustrates the resulting distribution of this loss. As seen in these graphs, our method manages to preserve the identity better than StarGAN in both datasets.

**Comparison to existing arts:** Our method obtains a better score than all three compared methods, StarGAN[3], Ganimation[24] and Gannotation[26], in every metric. This verifies that depth and surface normal information captures expression-specific details which are missed by sparse geometrical information, i.e. landmarks and action units.

**Hyper-parameter Study.** The important additional hyper-



Method	Geometry	Conf. penalty	Accuracy (%) $\uparrow$	FID $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$
StarGAN [3]	N/A	No	82.1	16.33	32.96	0.79
Ganination [24]	Action Units	No	74.3	18.86	32.12	0.74
Gannotation [26]	Landmarks	No	82.6	17.36	32.54	0.77
Our Method	Depth map	No	85.9(+3.8)	14.03	33.88	0.82
Our Method	Normal map	Yes	83.9(+1.8)	13.29	33.59	0.81
Our Method	Depth map	Yes	<b>87.2(+5.1)</b>	<b>13.35</b>	<b>35.29</b>	<b>0.84</b>

Table 1. Performance Comparison on AffectNet.

Method	Geometry	Conf. penalty	Accuracy (%) $\uparrow$	FID $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$
StarGAN [3]	N/A	No	53.1	39.45	31.19	0.75
Our Method	Depth map	Yes	<b>62(+8.9)</b>	<b>35.14(-4.31)</b>	<b>32.57(+1.38)</b>	<b>0.79(+0.04)</b>

Table 2. Performance Comparison on RaFD.

Method	Depth network weight ( $\beta$ )	Accuracy (%) $\uparrow$
StarGAN	0.0	82.1
Our Method	0.1	85.9(+3.8)
Our Method	0.2	86.1(+4.0)
Our Method	0.3	<b>86.3(+4.2)</b>
Our Method	0.4	86.3(+4.2)

Table 3. Performance Comparison on AffectNet with Different Weights for the Depth Network.




Figure 4. ID Loss Distribution. AffectNet (left) and RaFD (right). Please zoom in to see the details.

parameter for us in comparison to StarGAN is the weight of the loss incurred by the depth network, which comprises depth adversarial loss and depth classification loss. We set different weights for this loss and evaluate the performance. Table 3 summarises this hyper-parameter study. When we set the weight of depth loss to 0.0, which is the RGB Network alone, i.e. StarGAN, the mean target expression classification accuracy is 82.1%. As discussed before, without using the confidence penalty, setting the weight to 0.1 improves StarGAN by +3.8, yielding an accuracy rate of 85.9. We set the weight to 0.2, which improves the performance from 85.9% to 86.1%. Slightly increasing the weight to 0.3 gives a mean accuracy of 86.3%. As increasing the weight to 0.4 does not further improve the performance, we choose 0.3 as the optimal value for this parameter. These experimental results show that the performance of our method is influenced by the hyper-parameters but is stable.

## 4.2. Qualitative Evaluation


Figure 5 shows a comparison of some samples synthesised by StarGAN and our method. We observe that in-

corporating the depth information with our method yields results that are of higher quality while also maintaining the facial geometry and photometry, hence the identity.

On the left-hand side of Figure 5, we present some of the examples that show our method is capable of providing photometric consistency. Images synthesised by StarGAN contain artifacts on the skin, and in some cases fail to preserve the shapes of nose, mouth and eyes as well as the original texture. On the other hand, in almost all synthesised images our method yields results that look more realistic and natural. In particular, areas like mouth, nose and eyebrows are well synthesised in correspondence with the given target expression even in cases where StarGAN fails.

The right-hand side of Figure 5 presents a comparison of samples generated by StarGAN and our method along with their depth maps predicted by our method. For the first two and the last sample, when translating a neutral image to happy, our method synthesises the face with a realistic geometry whereas StarGAN fails in the mouth area resulting in artifacts in the inner mouth. In translation from happy to neutral on the fourth and fifth samples, StarGAN yields results where the mouth is uncertain and not closed. Our method, on the other hand, results in a closed mouth with little or no artifacts. The predicted depth maps for StarGAN verify this inference as the mouth is ambiguous in the fourth depth map and open in the fifth. For the third and sixth samples, when synthesising images with the target class 'sad', both methods generate open-mouth samples. However, in images generated by StarGAN the inner mouth has artifacts, the mouth is blurry and the original mouth geometry is not preserved, which is reflected in the predicted depth maps. These results show that the guidance of the depth network leads the RGB network to synthesise images that yield more realistic depth map predictions, which results in improved geometric and photometric consistency in the synthesised images. These qualitative examples further validate that the geometric consistency loss is essential while training GANs for high quality expressions translation.

Figure 6 shows a comparison of our method to existing



(a) Photometric consistency

(b) Geometric consistency

Figure 5. **Qualitative Results on AffectNet.** (a) A comparison of images synthesised by StarGAN and our method. Results show that our method preserves photometric consistency. (b) Synthesised images and their predicted depth maps. The original images with the given expressions are translated to the target expressions by StarGAN and our method. Our depth map estimator network is applied to obtain the predicted depth maps.

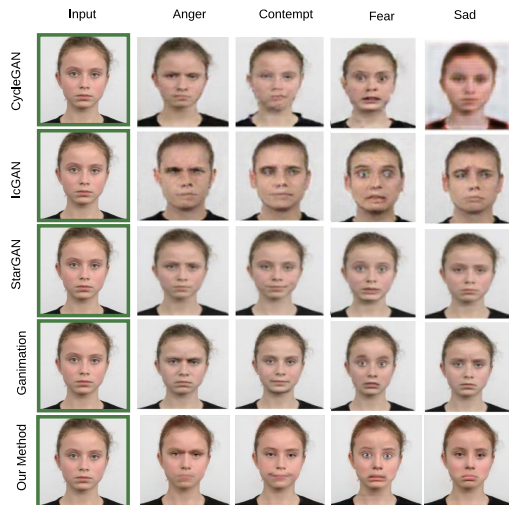


Figure 6. **Qualitative Results on RaFD.** Facial expressions synthesised by CycleGAN[35], IcGAN[22], StarGAN[3], GANimation[24] and our method. Synthesised images for previous methods are taken from [24].

work. In this figure we observe that our method produces outperforming or competitive results on RaFD as well. Results produced by CycleGAN and IcGAN are of lower quality, whilst StarGAN and GANimation produce high quality results. However, we observe that the geometric guidance introduced by our method helps preserving the existing geometry while modifying expression-specific geometrical

details. For instance, contempt expression synthesised by GANimation only slightly modifies the mouth, whereas our method adds mouth wrinkles and lowers the eyelids.

## 5. Conclusion

In this paper, we proposed a novel end-to-end deep network for manipulating facial expressions. The proposed method incorporates the depth information with an additional network that guides the training process by the depth consistency loss. To train the proposed pipeline of networks, we constructed a dataset that consists of RGB images and their corresponding depth maps and surface normal maps using an off-the-shelf 3D reconstruction model. To improve generalisation and lower the bias of the depth parameters, a confidence regulariser is applied to the discriminator side of the GAN frameworks. We evaluated our method on two challenging benchmarks, AffectNet and RaFD, and showed that our method synthesised samples that are both qualitatively and quantitatively superior to the ones generated by recent methods.

## Acknowledgements

R. Bodur is funded by the Turkish Ministry of National Education. This work is partly supported by EPSRC Programme Grant ‘FACER2VM’(EP/N007743/1). We thank Google Cloud Platforms for research credits award.



## References

- [1] Binod Bhattarai and Tae-Kyun Kim. Inducing optimal attribute representations for conditional gans. In *ECCV*, 2020.
- [2] Yuhua Chen, Wen Li, Xiaoran Chen, and Luc Van Gool. Learning semantic segmentation from synthetic data: A geometrically guided input-output adaptation approach. In *CVPR*, June 2019.
- [3] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *CVPR*, 2018.
- [4] Xing Di, Vishwanath A. Sindagi, and Vishal M. Patel. Gp-gan: Gender preserving gan for synthesizing faces from landmarks. In *ICPR*, 2018.
- [5] E Friesen and Paul Ekman. Facial action coding system: a technique for the measurement of facial movement. In *Consulting Psychologists Press*, 1978.
- [6] Baris Gecer, Binod Bhattarai, Josef Kittler, and Tae-Kyun Kim. Semi-supervised adversarial learning to generate photorealistic face images of new identities from 3d morphable model. In *ECCV*, 2018.
- [7] Rick Groenendijk, Sezer Karaoglu, Theo Gevers, and Thomas Mensink. On the benefit of adversarial training for monocular depth estimation. *CVIU*, 2020.
- [8] Zhenliang He, Wangmeng Zuo, Meina Kan, Shiguang Shan, and Xilin Chen. Arbitrary facial attribute editing: Only change what you want. *arXiv preprint arXiv:1711.10678*, 2017.
- [9] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017.
- [10] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS*, 2015.
- [11] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017.
- [12] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016.
- [13] Jean Kossaifi, Linh Tran, Yannis Panagakis, and Maja Pantic. Gagan: Geometry-aware generative adversarial networks. In *CVPR*, 2018.
- [14] Oliver Langner, Ron Dotsch, Gijsbert Bijlstra, Daniel Wigboldus, Skyler Hawk, and Ad Knippenberg. Presentation and validation of the radboud face database. *Cognition & Emotion*, 2010.
- [15] Mu Li, Wangmeng Zuo, and David Zhang. Deep identity-aware transfer of facial attributes. In *CoRR*, 2016.
- [16] Ming Liu, Yukang Ding, Min Xia, Xiao Liu, Errui Ding, Wangmeng Zuo, and Shilei Wen. Stgan: A unified selective transfer network for arbitrary image attribute editing. In *CVPR*, 2019.
- [17] Rui Min, Neslihan Kose, and Jean-Luc Dugelay. Kinectfacedb: A kinect database for face recognition. *Systems, Man, and Cybernetics: Systems, IEEE Transactions on*, 2014.
- [18] Ali Mollahosseini, Behzad Hasani, and Mohammad H. Mahoor. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing*, 10:18–31, 2019.
- [19] Koki Nagano, Jaewoo Seo, Jun Xing, Lingyu Wei, Zimo Li, Shunsuke Saito, Aviral Agarwal, Jens Fursund, and Hao Li. pagan: real-time avatars using dynamic textures. In *SIGGRAPH Asia 2018 Technical Papers*, page 258. ACM, 2018.
- [20] Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. In *BMVC*, 2015.
- [21] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NeurIPS*, 2017.
- [22] Guim Perarnau, Joost Van De Weijer, Bogdan Raducanu, and Jose M Álvarez. Invertible conditional gans for image editing. *arXiv preprint arXiv:1611.06355*, 2016.
- [23] Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. Regularizing neural networks by penalizing confident output distributions. *arXiv preprint arXiv:1701.06548*, 2017.
- [24] Albert Pumarola, Antonio Agudo, Aleix M Martinez, Alberto Sanfeliu, and Francesc Moreno-Noguer. Ganimation: Anatomically-aware facial animation from a single image. In *ECCV*, 2018.
- [25] Fengchun Qiao, Naiming Yao, Zirui Jiao, Zhihao Li, Hui Chen, and Harrison Wang. Geometry-contrastive gan for facial expression transfer. In *CoRR*, 2018.
- [26] Enrique Sanchez and Michel Valstar. Triple consistency loss for pairing distributions in gan-based face synthesis. *arXiv preprint arXiv:1811.03492*, 2018.
- [27] Klaus R Scherer. Emotion as a process: Function, origin and regulation, 1982.
- [28] Zhixin Shu, Ersin Yumer, Sunil Hadap, Kalyan Sunkavalli, Eli Shechtman, and Dimitris Samaras. Neural face editing with intrinsic image disentangling. In *Proceedings of the IEEE conference on*

*computer vision and pattern recognition*, pages 5541–5550, 2017.

- [29] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *ICLR*, 2015.
- [30] Lingxiao Song, Zhihe Lu, Ran He, Zhenan Sun, and Tieniu Tan. Geometry guided adversarial facial expression synthesis. In *ACM Multimedia*, 2018.
- [31] Anh Tuan Tran, Tal Hassner, Iacopo Masi, Eran Paz, Yuval Nirkin, and Gérard Medioni. Extreme 3d face reconstruction: seeing through occlusions. In *CVPR*, 2018.
- [32] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *Ieee Transactions on Image Processing*, 13, 2004.
- [33] Taihong Xiao, Jiapeng Hong, and Jinwen Ma. Elegant: Exchanging latent encodings with gan for transferring multiple face attributes. In *ECCV*, 2018.
- [34] Lijun Yin, Xiaozhou Wei, Yi Sun, Jun Wang, and Matthew J. Rosato. A 3d facial expression database for facial behavior research. In *FG*, 2006.
- [35] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.