

Synthetic Expressions are Better Than Real for Learning to Detect Facial Actions

Koichiro Niinuma¹, Itir Onal Ertugrul², Jeffrey F Cohn³, László A Jeni⁴

¹Fujitsu Laboratories of America, Pittsburgh, PA, USA, ²Tilburg University, Tilburg, The Netherlands

³University of Pittsburgh, Pittsburgh, PA, USA, ⁴Carnegie Mellon University, Pittsburgh, PA, USA

kniinuma@fujitsu.com, i.onal@uvt.nl, jeffcohn@pitt.edu, laszlojeni@cmu.edu

Abstract

Critical obstacles in training classifiers to detect facial actions are the limited sizes of annotated video databases and the relatively low frequencies of occurrence of many actions. To address these problems, we propose an approach that makes use of facial expression generation. Our approach reconstructs the 3D shape of the face from each video frame, aligns the 3D mesh to a canonical view, and then trains a GAN-based network to synthesize novel images with facial action units of interest. To evaluate this approach, a deep neural network was trained on two separate datasets: One network was trained on video of synthesized facial expressions generated from FERA17; the other network was trained on unaltered video from the same database. Both networks used the same train and validation partitions and were tested on the test partition of actual video from FERA17. The network trained on synthesized facial expressions outperformed the one trained on actual facial expressions and surpassed current state-of-the-art approaches.

1. Introduction

Facial expressions convey emotional states, behavioral intentions, and physical state [35]. In behavior sciences the gold-standard to decode such facial expressions is the Facial Action Coding System (FACS) [10]. FACS decomposes facial expressions into anatomically based action units (AUs), which alone or in combinations can represent nearly all possible facial expressions. While much progress has been made in action unit detection, at least two significant problems impede further advances.

First, while much video of spontaneous facial expression has been collected, only a fraction have been manually annotated. Expert FACS annotators are relatively few and the time required to comprehensively annotate action units slows the effort. AU annotation of a single minute of video typically requires one to three hours from a highly trained expert [8]. Reaching human-like accuracy in a fully super-

vised way would require labeled datasets orders of magnitude larger than those available today.

Second, AU labels are highly skewed in spontaneous behavior. Many AUs occur rarely and only a sparse subset of AU intensities occur at a time. As a consequence, rare classes do not contribute equally during classifier training, which hinders learning and undermines global performance. Although imbalanced learning has been well explored in the past, most approaches deal with a single majority and minority class and are not directly applicable to the multi-label AU domain.

We propose a generative semi-supervised method that can handle within a common framework both the limited sizes of annotated data now available and the low frequencies of occurrence.

Our approach makes use of a 3D facial expression generator trained on the labeled portion. The approach first reconstructs the 3D shape of the face from each video frame, aligns the reconstructed meshes to a canonical view to establish semantic correspondence across frames and subjects, and then trains a GAN-based network to synthesize novel images with facial action units of interest. We then use the network to generate a large AU-balanced dataset from unlabeled images for training.

Intensity of facial actions may be one of the most important features in assessing a person's emotional state [26]. Low intensity actions are detectable through motion [2]. Table 1 compares AU intensity datasets. In order to be able to detect these fine scale changes, we selected FERA 2017 [36], DISFA [25] and UNBC Pain [24], which are video datasets having manual fine-grained AU intensity annotations, to evaluate our approach (see supplementary materials for results on UNBC Pain). Note that Table 1 does not include datasets without AU intensity annotation, such as Aff-Wild2 [20].

Our novelties are twofold:

3D geometry based AU manipulation. Unlike previous work on facial AU manipulation that is limited to either 2D representations [30] or individual frames [13], our approach uses the 3D structure of the face to create semantic

Table 1. Comparison of AU intensity datasets

| | # of AUs with intensity codes | Continuous Video | Manual Ground Truth | Social Context | Manual Coding Reliability |
|-----------------------------|----------------------------------|---------------------|-------------------------------|-------------------|------------------------------|
| FERA 2017 [36] ⁱ | 7 | ✓ | ✓ | ✓ | Good |
| DISFA [25] | 12 | ✓ | ✓ | | Good |
| EmotioNet [5] | 12 ⁱⁱ | | Semi automated ⁱⁱⁱ | | Unknown |
| UNBC Pain [24] | 10 | ✓ | ✓ | | Good |

ⁱ The FERA 2017 dataset consists of BP4D [39] and BP4D+ [42].

ⁱⁱ In the EmotioNet Challenge 2020, 11 more AUs have been added for a total of 23 AUs. <http://cbcs1.ece.ohio-state.edu/enc-2020/>

ⁱⁱⁱ They manually FACS-coded 10% of this database.

correspondence across video-frames and subjects.

Synthetic multi-label stratification of AUs. Many AU occur infrequently, which undermines learning. To avoid imbalanced classes, we increase the prevalence and variety of under-represented AUs by synthesizing new facial expression.

2. Related Work

Solutions to limited AU annotations: While massive amount of facial expression data is available, high quality annotations of AU intensity labels are limited. To mitigate the problems in the AU annotations, weakly-supervised, semi-supervised, and self-supervised approaches have been proposed. Weakly-supervised approaches aim to exploit incomplete, inaccurate or inexact annotations to provide supervision. Zhao et al. [44] proposed a weakly supervised clustering approach utilizing a large set of web images with inaccurate annotations. The annotations were obtained from either pretrained models or query strings. Ruiz et al. [32] proposed to train AU detectors without any AU annotations by leveraging the expression labels and using prior knowledge on expression-dependent AU probabilities. Similarly, Zhang et al. [41] exploited expression-dependent and expression-independent joint AU probabilities as prior knowledge and learned to detect AUs without any AU annotation. In another study, Zhang et al. [40] used various types of domain knowledge including relative appearance similarity, temporal intensity ordering, facial symmetry, and contrastive appearance difference to provide weak supervision for AU intensity estimation with extremely limited annotations. Peng et al. [29] proposed a method that learns AU classifiers from domain knowledge and expression-annotated facial images through adversarial training.

Semi-supervised approaches deal with partially annotated data. They aim to leverage the unlabelled data with the assumption that unlabelled data follow continuity or form cluster with the labeled data [44]. Wu et al. [37] used Restricted Boltzmann Machine to model the AU distribution using the annotated labels, which is used to train the AU classifiers with partially labeled data. Zeng et al. [38] trained a quasi-semi-supervised (QSS) classifier with

virtual labels provided by the confident positive and negative classifiers, which separate easily identified positive and negative samples from all else, respectively. Niu et al. [28] proposed a semi-supervised co-training approach named as multi-label coregularization for AU recognition, which aims to improve AU recognition with abundant unlabelled face images and domain knowledge of AUs.

Recent works focus on self-supervised approaches where the goal is to learn the discriminative representation from the massive amount of videos without annotations. Li et al. [22] proposed a self-supervised learning framework named Twin-Cycle Autoencoder that disentangles the AU-related movements from the pose-related ones to learn AU representations from unlabelled videos. While the aforementioned approaches tackle the problems in the annotations, none of them aims to balance the distribution of AU intensities. For label balancing, upsampling approaches where the infrequent labels are selected multiple times [21][43] or multi-label minority oversampling majority undersampling approach [7] have been used. Since resampling is done within the dataset, such balancing methods do not contribute additional semantic information about the infrequent label.

GAN-based facial expression transfer: Recently GANs have received attention to transfer facial expressions from a source subject to a target subject. Existing work on GAN-based facial expression transfer approaches focus on generating facial images with discrete emotions [6][9], or the specified facial action units [30][23]. Some of the GAN-based approaches specifically aim to guide their models with the facial geometry information. Song et al. [34] proposed a Geometry-Guided Generative Adversarial Network (G2-GAN) which employs fiducial points as a controllable condition to guide facial texture synthesis with specific expression. Qiao et al. [31] applied contrastive learning in GAN to embed geometry information onto a semantic manifold in the latent space for facial expression transfer. Geng et al. [13] combined 3DMMs and deep generative techniques in a single framework for fine-grained face manipulation. Yet, in these studies transfer was limited to either 2D representation or to individual frames.

Synthetic data augmentation: Some of the recently

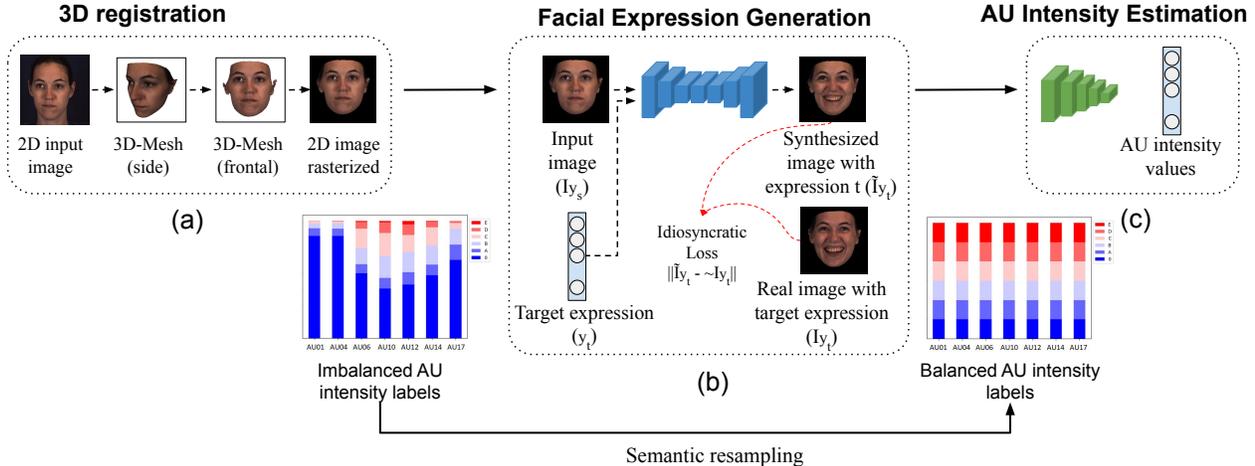


Figure 1. Pipeline of our approach. (a) A dense 3D mesh of the face is reconstructed, rotated to the canonical frontal view and 2D rasterized. (b) Given the 3D registered input image (\mathbf{I}_{y_s}) and target expression (\mathbf{y}_t), target expression is transferred to input image ($\tilde{\mathbf{I}}_{y_t}$). Using semantic resampling on the synthesized images, a training set of balanced AU intensity labels is formed. (c) Obtained balanced training set is used to train a convolutional neural network to estimate AU intensity values.

proposed methods utilized synthetic data for facial expression analysis. Abbasnejad et al. [1] pre-trained their model using synthetic face images and then fine-tuned it on real images. Zhu et al. [46] proposed a data augmentation method using GAN to classify basic emotions. Kollias et al. [19] proposed an approach using 3DMMs to synthesize facial affect: in terms of six basic emotions or in terms of valence and arousal. Unlike the existing methods, our approach is designed to generate a large AU-balanced dataset.

3. Proposed Method

Fig. 1 shows the pipeline of our approach. First we perform dense 3D registration from 2D images. Then, we train a GANimation-based architecture with idiosyncratic loss to synthesize new facial expressions. Using semantic resampling, we obtain a balanced distribution of AU intensity labels. Finally, using the balanced synthetic expression database, we train a convolutional neural network architecture for AU intensity estimation.

3.1. 3D face registration

We normalize videos using PRNet [12], a face alignment software that accomplishes dense 3D registration from 2D images without requiring person-specific training. PRNet uses an encoder-decoder architecture containing convolutional layers and residual blocks to jointly perform facial landmark alignment and 3D facial structure reconstruction. This architecture learns a mapping from an RGB image to UV position map (a 2D image representation of 3D coordinates in UV space keeping the position and semantic information). By learning the position map, it is possible to

directly regress the complete 3D structure along with semantic meaning from a single image. Using PRNet, we obtain the dense 3D mesh of the face in a frontal view and texture information. Then, we map the texture to 3D mesh and rasterize it to 2D image of size 224 x 224.

3.2. Facial Expression Generation Architecture

We build upon GANimation [30] framework to synthesize novel facial expressions. First we map AU intensity labels (0 to E-level) to values in range [0,1]. Given a 3D registered source image \mathbf{I}_{y_s} with the AU intensity values $\mathbf{y}_s = \{s_1, s_2, \dots, s_n\}$, and target AU intensity values $\mathbf{y}_t = \{t_1, t_2, \dots, t_n\}$, we synthesize $\tilde{\mathbf{I}}_{y_t}$. With our architecture, we aim to minimize these following terms:

Image Adversarial Loss: In order to obtain realistic synthesized images and ensure that the distribution of the generated images are similar to the distribution of the training images, we use image adversarial loss. Let \mathcal{P}_s be the data distribution of the source image, $\mathcal{P}_{\tilde{I}}$ be the random interpolation distribution, and λ_{gp} be the penalty loss. Then we can write the image adversarial loss $\mathcal{L}_{adv}(G, D_{adv}, \mathbf{I}_{y_s}, \mathbf{y}_t)$ as follows:

$$\mathbb{E}_{\mathbf{I}_{y_s} \sim \mathcal{P}_s} [D_{adv}(G(\mathbf{I}_{y_s} | \mathbf{y}_t))] - \mathbb{E}_{\mathbf{I}_{y_s} \sim \mathcal{P}_s} [D_{adv}(\mathbf{I}_{y_s})] + \lambda_{gp} \mathbb{E}_{\tilde{I} \sim \mathcal{P}_{\tilde{I}}} [(\|\nabla_{\tilde{I}} D_{adv}(\tilde{I})\|_2 - 1)^2] \quad (1)$$

where G denotes generator and D_{adv} denotes adversarial discriminator.

Conditional Expression Loss: In order to enforce G to synthesize images containing the target expression \mathbf{y}_t , we use the following loss $\mathcal{L}_{exp}(G, D_{exp}, \mathbf{I}_{y_s}, \mathbf{y}_s, \mathbf{y}_t)$ to mini-

mize the distance between AU intensities of the source and target images:

$$\mathbb{E}_{\mathbf{I}_{y_s} \sim \mathcal{P}_s} [\|D_{exp}(G(\mathbf{I}_{y_s} | \mathbf{y}_t)) - \mathbf{y}_t\|_2^2] + \mathbb{E}_{\mathbf{I}_{y_s} \sim \mathcal{P}_s} [\|D_{exp}(\mathbf{I}_{y_s}) - \mathbf{y}_s\|_2^2] \quad (2)$$

where D_{exp} denotes discriminator for expression.

Identity Loss: We aim to guarantee that the face in both the input and output images belong to the same person. We use this cycle-consistency loss to penalize the difference between the original image \mathbf{I}_{y_s} and its reconstruction \mathbf{I}_{y_t} .

$$\mathcal{L}_{idt}(G, \mathbf{I}_{y_s}, y_s, y_t) = \mathbb{E}_{\mathbf{I}_{y_s} \sim \mathcal{P}_s} [\|G(G(\mathbf{I}_{y_s} | y_t) | y_s) - \mathbf{I}_{y_s}\|_1] \quad (3)$$

Idiosyncratic Loss: With the GANimation architecture, we can transfer the AU intensity values \mathbf{y}_t of a target image \mathbf{J}_{y_t} to the source image \mathbf{I}_{y_s} to synthesize $\tilde{\mathbf{I}}_{y_s}$. When the identity of source (I) and target (J) images are the same, then we can minimize the difference between \mathbf{I}_{y_t} and $\tilde{\mathbf{I}}_{y_t}$ to ensure that both expression and identity of the synthesized image are the same as the target image. Idiosyncratic loss can be defined as:

$$\mathcal{L}_{ids}(G, \mathbf{I}_{y_s}, y_s, y_t) = \mathbb{E}_{\mathbf{I}_{y_s} \sim \mathcal{P}_s} [\|(G(\mathbf{I}_{y_s} | y_t) - \mathbf{I}_{y_t})\|_1] \quad (4)$$

Final Loss: We obtain our final loss by combining of the mentioned individual losses as follows:

$$\mathcal{L} = \lambda_{adv} \mathcal{L}_{adv} + \lambda_{exp} \mathcal{L}_{exp} + \lambda_{idt} \mathcal{L}_{idt} + \lambda_{ids} \mathcal{L}_{ids} \quad (5)$$

where λ_{adv} , λ_{exp} , λ_{idt} , and λ_{ids} are the hyperparameters used to adjust the importance of different components.

3.3. AU intensity estimation

After we synthesize new expressions for the individuals in our database, we perform semantic resampling and create a training set having balanced AU intensity labels for each AU. Then we train convolutional neural networks (VGG16) using the balanced synthetic training set. During test time, we obtain AU intensity outputs of each estimator.

4. Experiments

4.1. Datasets

In all of our experiments, we used three facial expression datasets. For training the generator and evaluating within domain performance, we used the widely accepted 2017 Facial Expression Recognition Benchmark (FERA 2017) [36]. For generating out-of-domain samples, we used the high resolution images from MultiPIE [14]. To evaluate the generalizability of our AU classifiers to another domain,

we used the Denver Intensity of Spontaneous Facial Action Database (DISFA) [25].

FERA 2017: The FERA 2017 Challenge was the first to provide a common protocol with which to compare approaches to detection of AU occurrence and AU intensity robust to pose variation. FERA 2017 provided synthesized face images with 9 head poses as shown in Fig.5. The training set is based on the BP4D database [39], which includes digital videos of 41 participants. The development and test sets are derived from BP4D+ [42] and include digital videos of 20 and 30 participants, respectively. FERA 2017 presented two sub-challenges: occurrence detection and intensity estimation. For the former 10 AUs were labelled; for the latter, 7 AUs were labelled. For our experiments, 7 AUs for intensity estimation were used.

MultiPIE: The MultiPIE dataset contains images of 337 people recorded in up to four sessions over the span of five months. Subjects were imaged under 15 view points and 19 illumination conditions while displaying a range of facial expressions. In addition, high resolution frontal images were acquired as well. For synthesis, we used the high resolution images only.

DISFA: The DISFA dataset contains videos of 27 adult subjects (12 women, 15 men). It is manually annotated for AU intensity from 0 to E-level. Participants watched a video clip consisting of 9 segments intended to elicit a range of facial expressions of emotion.

4.2. Experimental Setup

In this section we describe the experimental setup for the generator network and the classifier.

4.2.1 Generator Training

In our experiments, we used the FERA 2017 dataset [36] to train facial expression generation models. The dataset consists of Train, Valid, and Test partition. We used Train partition only to train our model. All the images were resized to 224×224 pixels to match the receptive field of our AU estimation model (VGG16).

In all of our experiments, we used a GANimation [30] replicate implementation¹. We modified the loss function with the idiosyncratic constraint as described in the previous section.

4.2.2 Classifier Training

For the baseline experiments, we used the training partition of the FERA 2017 dataset [36] to train AU classifiers and

¹https://github.com/donydchen/ganimation_replicate

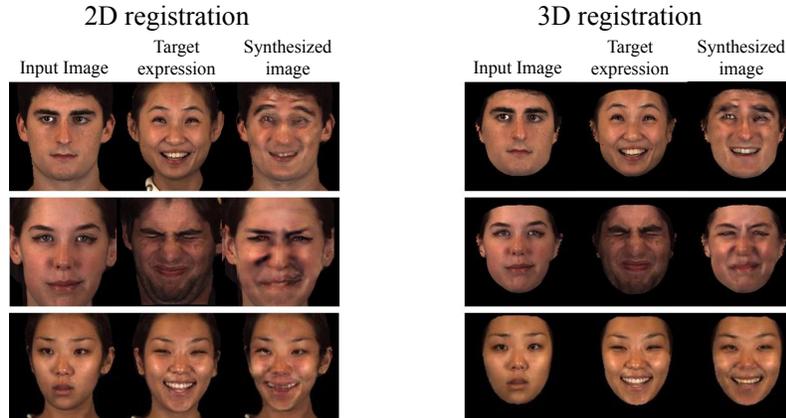


Figure 2. Examples of 2D and 3D normalization. Note the higher image quality when 3D normalization was used.

used the test partition to test them. To create a balanced training set and compare methods, 5,000 frames were selected for each AU intensity. In the case of real images, we randomly selected 5,000 images from the six intensity classes (not present, and A to E levels). We down-sampled the majority classes and up-sampled the minority classes to reach this number. In the case of synthesized images, we first selected 5,000 pairs of input images and target AU labels for each intensity and each AU. Then we synthesized 5,000 images for these pairs. We selected 5,000 images according to experimental results by Niinuma et al. [27]. They analyzed the influence of training set size on FERA17, and showed that the training set sized have minor influence on the performance: score peaked at 5,000 images, after that performance plateaued. During the selection input images were selected only from frames that did not have the target AU of interest. For example, when we synthesized images for E level intensity of AU1, we randomly selected 5,000 input images that did not have AU1 present and 5,000 target AU labels having E level intensity for AU1. We employ this strategy for two main reasons. First, while the generator can add realistic facial deformations (like wrinkles and bulges) to neutral faces, it oftentimes fails in removing those. Therefore, to acquire higher quality synthesized images, starting from neutral frame is preferable. Second, since the AU labels are sparse, there are many more frames in the dataset where the target AU is not present. This way we can obtain a variety of synthesized images for each facial expression.

We selected a VGG16 network pre-trained on ImageNet for the baseline architecture for AU estimation. Previous studies found this combination preferable for AU coding [27]. We replaced the final layer of the network with a 6-length one-hot representation, and fine-tuned VGG16 network from the third convolutional layer. Dropout rate was set to 0.5, and Adam optimizer was used with $LR =$

5×10^{-5} as suggested in [27].

4.3. Synthetic vs Real Expressions under 2D vs 3D Alignment

In this set of experiments we studied how two main components affect the performance of the whole system.

First, we were interested in the effect of face alignment on the synthesis and AU recognition performance. We explored both 2D and 3D alignment. 2D normalization treats the face as 2D object. That assumption is reasonable, as long as there is no head movement present. As soon as head orientation deviates from frontal, one expects the classifier’s ability to measure expressions to degrade. On the other hand, 3D normalization should be able to preserve semantic correspondences of the different facial regions across poses, and result in higher performance. For 2D alignment, we applied Procrustes analysis between 68 landmarks provided by the dlib face tracker[18] on the frames and a frontal template. For 3D normalization we used the method described in Sec. 3.1. Note that there are not significant differences between shallow and deep approaches in terms of facial alignment [16][33]. In the 300-W Challenge [33], the two top methods are a cascade regressor and a CNN approaches. The dlib face tracker is an implementation of a decision tree based cascade regressor [17].

Second, we were interested how synthetic expressions would affect the classification performance. We compared multi-label minority oversampling & majority undersampling with the proposed, completely synthetic expression generation. Both methods balance the skewed distributions of AUs, but while the first one can not produce more varied minority samples, the latter one can. As mentioned in Sec. 4.2.2, we used 5,000 images for each intensity each AU for both real and synthesized image settings.

Fig. 2 shows examples of synthetic expressions generating using 2D and 3D alignment. 3D alignment results in less

Table 2. Comparison of synthetic vs real expressions under 2D vs 3D alignment. Scores are Inter-rater reliability (ICC) of AU intensity level estimation. The Images row shows which is used to train classifiers: Real or Synthetic. All of the classifiers were tested on the real test dataset. The same Registration (2D or 3D) were applied to both train and test datasets.

| Registration Images | 2D | | 3D | |
|---------------------|--------------|--------------|--------------|--------------|
| | Real | Synthetic | Real | Synthetic |
| AU1 | 0.431 | 0.336 | 0.343 | 0.381 |
| AU4 | 0.223 | 0.116 | 0.260 | 0.219 |
| AU6 | 0.796 | 0.790 | 0.751 | 0.804 |
| AU10 | 0.777 | 0.812 | 0.785 | 0.773 |
| AU12 | 0.801 | 0.792 | 0.806 | 0.795 |
| AU14 | 0.118 | 0.238 | 0.084 | 0.244 |
| AU17 | 0.395 | 0.374 | 0.391 | 0.461 |
| Mean | 0.506 | 0.494 | 0.489 | 0.525 |

ghosting and other texture artifacts and provide higher quality images. Inter-rater reliability (ICC) results of AU classification under these four different conditions are shown in Table. 2. 3D registration with synthetic re-sampling outperformed the other three conditions.

In these experiments we conducted a parameter search to find the optimal values of the generator for classification. We varied parameters that control the contribution of the adversarial loss (λ_{adv}), conditional expression loss (λ_{exp}), and the identity loss (λ_{idt}). In our baseline configuration, we used 1.0 for λ_{adv} , 160 for λ_{exp} , and 10 for λ_{idt} . Fig. 3 shows the impact of parameter values on the intensity estimation for different AUs. λ_{idt} does not affect the classification significantly. For both λ_{exp} and λ_{adv} , we selected the global optimal values. Pumarola et al. introduced an attention mask and a color transformation term in the lost function to prevent the attention mask saturation [30]. In our experiments we did not observe this saturation effect and removed these terms from the loss function.

4.4. The Effect of Idiosyncratic Loss

3D alignment improves the performance, image quality is still low in some cases especially when target expression has a high AU intensity. To mitigate the problem, we introduced a new idiosyncratic term in the loss function. FERA 2017 datasets include many images with different facial expressions for each subject. Idiosyncratic loss utilizes this feature of the datasets. Fig. 4 shows some of the examples ($\lambda_{ids} = 1$). The ICC result with the new function is 0.523. To calculate the ICC result, 3D synthetic train images were used to train classifiers, and 3D real test images were used to test the classifiers. Compared with the ICC without it (0.525), it does not improve the ICC, but we confirmed that it improves the image quality. The Frchet Inception Distance (FID) [15] for the synthetic dataset with idiosyncratic loss (4.88) is better than the one without idiosyncratic loss (5.94). Note that a lower FID is better.

4.5. Temporal Normalization and Comparison with State of the Art

Results from the previous experiment suggest that precise spatial alignment improves the performance. AUs are temporal, we decided to test the best method under different temporal normalization. We compared two methods that enhances the temporal aspect of the AUs: AU0 normalization and mean texture normalization. AU0 normalization computes the appearance differences between the actual frame and a neutral frame. A neutral frame is not necessary available in real life conditions, but we can assume that we have multiple frames from a single person. In personal mean texture normalization we calculate the average appearance of a person and then calculate the differences between each frame and the mean texture. This step minimizes individual differences in the appearance space.

The last three columns of Table. 3 show the results. Personal mean texture normalization shows the best results. On average, there is no gain with AU0 normalization, however, individual AU level differences are significant. While AU0 normalization shows better results for AU1 and AU4, it shows worse results for AU14 and AU17.

Table. 3 shows the comparison with other state-of-the-art methods using only the frontal poses. We report Inter-rater reliability (ICC) that is the standard metric of the FERA 2017 benchmark. Our method (with or without temporal enhancement) outperforms all other methods. For a fair comparison, we compared our approach with existing methods' results on frontal view of the test partition.

4.6. Experiment with Non-frontal Poses

Encouraged by the results of the previous experiment, we decided to evaluate the method's ability to generalize to unseen poses. FERA 2017 has nine different poses, we report the performance on all of these using the Test partition. We investigated three scenarios:

- (a) **real 2D normalized** \rightarrow **real 2D normalized**. For a baseline, we used real images with 2D alignment for training, and evaluated performance on 2D normalized real images from the testing set.
- (b) **synthetic 3D normalized** \rightarrow **real 3D normalized** We trained on synthetic 3D normalized images, and tested on 3D normalized real images from the test set. We applied the 3D normalization procedure described in Sec. 3.1 to each test image with non-frontal pose. Self-occluded facial parts were filled with black color during the rasterization step.
- (c) **3D augmented synthetic** \rightarrow **real 2D normalized** We synthesized 3D meshes and rotated them into the nine standard orientations found in FERA 2017. We randomly selected 500 images for each intensity, AUs,

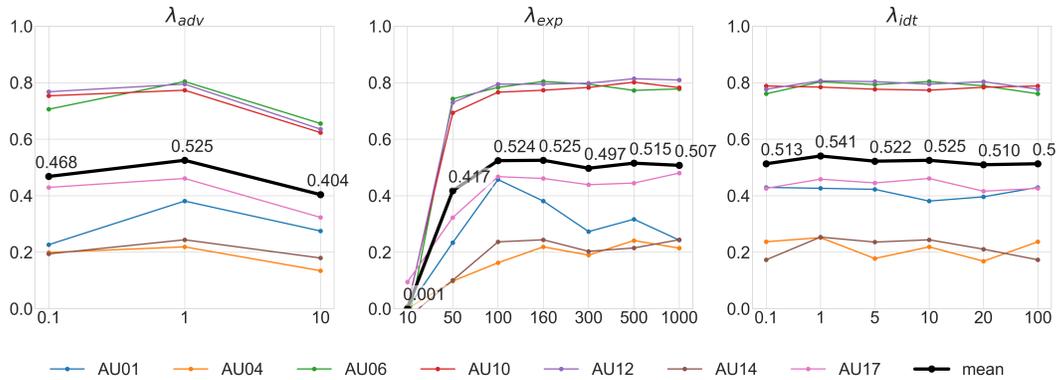


Figure 3. Influence of parameter values on ICC for intensity estimation



Figure 4. Comparison without and with idiosyncratic loss

Table 3. Comparison with state-of-the-art methods. Reported scores are Inter-rater reliability (ICC) on frontal views only. The results for No norm is the same with the ones for 3D Synthetic in Table 2.

| | Valstar | Amirian | Batista | Zhou | Niinuma | Ours | | |
|------|-------------|------------|------------|--------------|-------------|--------------|--------------|--------------|
| | et al. [36] | et al. [3] | et al. [4] | et al. [45] | et al. [27] | No norm | AU0 norm | Mean norm |
| AU1 | 0.025 | 0.270 | 0.311 | 0.286 | 0.433 | 0.381 | 0.539 | 0.613 |
| AU4 | 0.003 | 0.074 | 0.098 | 0.130 | 0.281 | 0.219 | 0.361 | 0.409 |
| AU6 | 0.616 | 0.644 | 0.721 | 0.625 | 0.786 | 0.804 | 0.764 | 0.779 |
| AU10 | 0.662 | 0.733 | 0.741 | 0.739 | 0.768 | 0.773 | 0.787 | 0.757 |
| AU12 | 0.709 | 0.745 | 0.754 | 0.822 | 0.812 | 0.795 | 0.792 | 0.794 |
| AU14 | 0.066 | 0.030 | 0.127 | 0.075 | 0.153 | 0.244 | 0.114 | 0.170 |
| AU17 | 0.015 | 0.271 | 0.252 | 0.342 | 0.382 | 0.461 | 0.288 | 0.465 |
| Mean | 0.299 | 0.395 | 0.429 | 0.431 | 0.516 | 0.525 | 0.521 | 0.570 |

poses. The total number of images for each intensity each AU is 4,500 while 5,000 images are selected for scenarios (a) and (b). We tested the system on 2D aligned test images.

Fig. 5 shows the results. While (a) and (b) show low ICCs when face poses are largely different from frontal views, the performance drop for (c) is much smaller. The re-

sults show that our synthesized images with 3D registration are also effective to non-frontal views by recreating non-frontal view images from the synthesized images. Note that: 1) the approaches in Table. 5 use frontal view images only to train models while the methods in FERA17 challenges used images with all 9 poses to train models., 2) the reason why the performance on frontal view for (c) is worse than

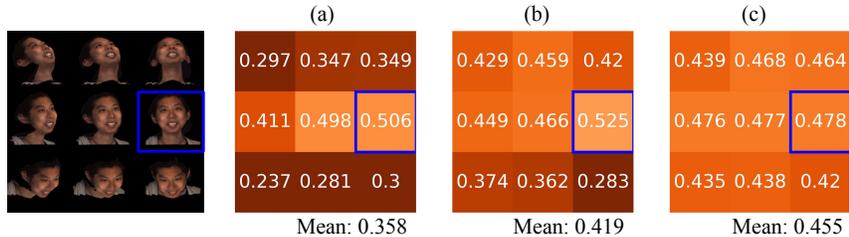


Figure 5. ICC for Test partition with non-frontal view.

Table 4. Cross domain ICC performance. (Synthetic training set \rightarrow Real test set)

| | MultiPIE \rightarrow FERA | FERA \rightarrow DISFA | MultiPIE \rightarrow DISFA |
|------|-----------------------------|--------------------------|------------------------------|
| AU01 | 0.311 | 0.314 | 0.418 |
| AU04 | 0.202 | 0.400 | 0.541 |
| AU06 | 0.786 | 0.573 | 0.524 |
| AU10 | 0.726 | - | - |
| AU12 | 0.792 | 0.748 | 0.698 |
| AU14 | 0.168 | - | - |
| AU17 | 0.365 | 0.373 | 0.290 |
| Mean | 0.479 | 0.482 | 0.494 |

(a) and (b) is that only 500 frontal view images for each intensity each AU are used to train.

4.7. Cross-domain Experiments

We have learned from the previous experiments that our AU classifiers can perform well when trained and tested within the same domain. To evaluate the generalizability of our approach to unseen domains (both in generating expressions and evaluating classifiers), we conducted two sets of experiments.

First, we were interested in how generating out-of-domain samples would affect the performance on the FERA 2017 Test partition. In this case we trained the generator network on FERA 2017 Train partition, but we synthesized new expressions using high resolution frontal images present in the MultiPIE dataset. We selected 921 MultiPIE images and generated 5 target expressions each, resulting in 4,605 images for each intensity and each AU. We trained the classifier on these images and tested the performance on FERA 2017 Test partition.

In the second cross-domain experiment we evaluated out-of-domain classification. Here classifiers were trained either on synthetic expressions generated from FERA 2017 or synthetic expressions generated from MultiPIE, and they were tested on DISFA. The DISFA dataset differs in imaging condition and type of AU coding: Context is not social in DISFA while in FERA subjects are interacting with the experimenter, and in DISFA, the base rates of most AUs is very low and limited to what occurs in a film-watching paradigm [11].

In all of these experiments, FERA17 Train partition was used to train facial expression generation models, and

3D normalization was applied to each image. The whole dataset of DISFA was used to test the models.

Table. 4 shows the results. Performance of models trained on synthesized MultiPIE expressions (0.479) is lower than the one trained on synthesized FERA 2017 expression (0.525), but there is only 1% difference with the one trained on real FERA17 expressions (0.489). The results on DISFA shows that the result trained on synthesized MultiPIE expressions (0.494) is slightly better than the one trained on synthesized FERA17 expressions (0.482).

5. Conclusion

We have proposed a generative approach that achieves 3D geometry based AU manipulation with idiosyncratic loss to synthesize facial expressions. With the semantic re-sampling, our approach provides a balanced distribution of AU intensity labels, which is crucial to train AU intensity estimators. We have shown that using the balanced synthetic set for training performs better than using the real training dataset on the same test set. Generating expressions using the 3D registered facial images gives better AU intensity estimation performance compared to using 2D registered ones. Moreover, our proposed idiosyncratic loss has improved the visual quality of the outputs. Cross-pose and cross-domain results reveal that classifiers trained on our synthesized images are also effective to non-frontal views and to unseen domains.

6. Acknowledgments

This research was supported in part by Fujitsu Laboratories, NIH awards NS100549 and MH096951, and NSF award CNS-1629716.

References

- [1] Iman Abbasnejad, Sridha Sridharan, Dung Nguyen, Simon Denman, Clinton Fookes, and Simon Lucey. Using synthetic data to improve facial expression analysis with 3d convolutional networks. *International Conference on Computer Vision Workshops (ICCVW)*, 2017.
- [2] Zara Ambadar, Jonathan W Schooler, and Jeffrey F Cohn. Deciphering the enigmatic face: The importance of facial dynamics in interpreting subtle facial expressions. *Psychological Science*, 2005.
- [3] M. Amirian, M. Kächele, G. Palm, and F. Schwenker. Support vector regression of sparse dictionary-based features for view-independent action unit intensity estimation. In *Automatic Face & Gesture Recognition (FG 2017), 2017 12th IEEE International Conference on*, pages 854–859, 2017.
- [4] J. C. Batista, V. Albiero, O. R. P. Bellon, and L. Silva. AUMPNNet: simultaneous action units detection and intensity estimation on multipose facial images using a single convolutional neural network. In *Automatic Face & Gesture Recognition (FG 2017), 2017 12th IEEE International Conference on*, pages 868–871, 2017.
- [5] C. F. Benitez-Quiroz, R. Srinivasan, and A. M. Martinez. Emotionet: An accurate, real-time algorithm for the automatic annotation of a million facial expressions in the wild. 2016.
- [6] Yunje 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [7] Wen-Sheng Chu, Fernando De la Torre, and Jeffrey F Cohn. Learning facial action units with spatiotemporal cues and multi-label sampling. *Image and vision computing*, 81:1–14, 2019.
- [8] Jeffrey F Cohn and Paul Ekman. Measuring facial action. *The new handbook of methods in nonverbal behavior research*, pages 9–64, 2005.
- [9] Hui Ding, Kumar Sricharan, and Rama Chellappa. Exprgan: Facial expression editing with controllable expression intensity. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [10] P Ekman, W Friesen, and J Hager. Facial action coding system: Research nexus. *Network Research Information, Salt Lake City, UT*, 1, 2002.
- [11] Itir Onal Ertugrul, Jeffrey F. Cohn, László A. Jeni, Zheng Zhang, Lijun Yin, and Qiang Ji. Crossing domains for au coding: Perspectives, approaches, and measures. *IEEE Transactions on Biometrics, Behavior, and Identity Science (TBIOM)*, 2020.
- [12] Yao Feng, Fan Wu, Xiaohu Shao, Yanfeng Wang, and Xi Zhou. Joint 3d face reconstruction and dense alignment with position map regression network. In *ECCV*, 2018.
- [13] Zhenglin Geng, Chen Cao, and Sergey Tulyakov. 3d guided fine-grained face manipulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9821–9830, 2019.
- [14] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker. Multi-pie. *Image and Vision Computing*, 28:807–813, 2010.
- [15] 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. *Neural Information Processing Systems (NIPS)*, 2017.
- [16] László A. Jeni, Sergey Tulyakov, Lijun Yin, Nicu Sebe, and Jeffrey F. Cohn. The first 3d face alignment in the wild (3dfaw) challenge. *European Conference on Computer Vision Workshops*, pages 511–520, 2016.
- [17] Vahid Kazemi and Josephine Sullivan. One millisecond face alignment with an ensemble of regression trees. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [18] D. E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10:1755–1758, 2009.
- [19] Dimitrios Kollias, Shiyang Cheng, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. Deep neural network augmentation: Generating faces for affect analysis. *International Journal of Computer Vision (IJCV)*, 2020.
- [20] Dimitrios Kollias and Stefanos Zafeiriou. Expression, affect, action unit recognition: Aff-wild2, multi-task learning and arcfac. *British Machine Vision Conference (BMVC)*, 2019.
- [21] Wei Li, Farnaz Abtahi, Zhigang Zhu, and Lijun Yin. Eac-net: A region-based deep enhancing and cropping approach for facial action unit detection. In *2017 12th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2017)*, pages 103–110. IEEE, 2017.
- [22] Y. Li, J. Zeng, S. Shan, and X. Chen. Self-supervised representation learning from videos for facial action unit detection. 2019.
- [23] Zhilei Liu, Diyi Liu, and Yunpeng Wu. Region based adversarial synthesis of facial action units. In *International Conference on Multimedia Modeling*, pages 514–526. Springer, 2020.
- [24] P. Lucey, J. F. Cohn, K. M. Prkachin, P. E. Solomon, and I. Matthews. Painful data: The unbc-mcmaster shoulder pain expression archive database. In *Automatic Face & Gesture Recognition and Workshops (FG 2011), 2011 IEEE International Conference on*, pages 57–64, 2011.
- [25] S. M. Mavadati, M. H. Mahoor, K. Bartlett, P. Trinh, and J. F. Cohn. DISFA : A spontaneous facial action intensity database. *IEEE Transactions on Affective Computing*, 4(2):151–160, 2013.
- [26] Gary McKeown, Ian Sneddon, and William Curran. Gender differences in the perceptions of genuine and simulated laughter and amused facial expressions. *Emotion Review*, 7(1):30–38, 2015.
- [27] K. Niinuma, L. A. Jeni, I. O. Ertugrul, and J. F. Cohn. Unmasking the devil in the details: What works for deep facial action coding? *British Machine Vision Conference (BMVC)*, pages 1–12, 2019.
- [28] Xuesong Niu, Hu Han, Shiguang Shan, and Xilin Chen. Multi-label co-regularization for semi-supervised facial action unit recognition. In *Advances in Neural Information Processing Systems*, pages 907–917, 2019.

- [29] G. Peng and S. Wang. Weakly supervised facial action unit recognition through adversarial training. pages 2188–2196, 2018.
- [30] A. Pumarola, A. Agudo, A.M. Martinez, A. Sanfeliu, and F. Moreno-Noguer. Ganimation: One-shot anatomically consistent facial animation. *International Journal of Computer Vision (IJCV)*, 2019.
- [31] Fengchun Qiao, Naiming Yao, Zirui Jiao, Zhihao Li, Hui Chen, and Hongan Wang. Geometry-contrastive gan for facial expression transfer. *arXiv preprint arXiv:1802.01822*, 2018.
- [32] Adria Ruiz, Joost Van de Weijer, and Xavier Binefa. From emotions to action units with hidden and semi-hidden-task learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3703–3711, 2015.
- [33] Christos Sagonas, Epameinondas Antonakos, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 300 faces in-the-wild challenge: database and results. *Image and Vision Computing*, 47:3–18, 2016.
- [34] Lingxiao Song, Zhihe Lu, Ran He, Zhenan Sun, and Tieniu Tan. Geometry guided adversarial facial expression synthesis. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 627–635, 2018.
- [35] Y-I Tian, Takeo Kanade, and Jeffrey F Cohn. Recognizing action units for facial expression analysis. *IEEE Transactions on pattern analysis and machine intelligence*, 23(2):97–115, 2001.
- [36] M. F. Valstar, E. Sánchez-Lozano, J. F. Cohn, L. A. Jeni, J. M. Girard, Z. Zhang, L. Yin, and M. Pantic. FERA 2017 - addressing head pose in the third facial expression recognition and analysis challenge. In *Automatic Face & Gesture Recognition (FG 2017), 2017 12th IEEE International Conference on*, pages 839–847, 2017.
- [37] Baoyuan Wu, Siwei Lyu, Bao-Gang Hu, and Qiang Ji. Multi-label learning with missing labels for image annotation and facial action unit recognition. *Pattern Recognition*, 48(7):2279–2289, 2015.
- [38] Jiabei Zeng, Wen-Sheng Chu, Fernando De la Torre, Jeffrey F Cohn, and Zhang Xiong. Confidence preserving machine for facial action unit detection. In *Proceedings of the IEEE international conference on computer vision*, pages 3622–3630, 2015.
- [39] X. Zhang, L. Yin, J. F. Cohn, S Canavan, M. Reale, A. Horowitz, and J. M. Girard. BP4D-spontaneous: a high-resolution spontaneous 3D dynamic facial expression database. *Image and Vision Computing*, 32(10):692–706, 2014.
- [40] Y. Zhang, W. Dong, B. Hu, and Q. Ji. Weakly-supervised deep convolutional neural network learning for facial action unit intensity estimation. pages 2314–2323, 2018.
- [41] Yong Zhang, Weiming Dong, Bao-Gang Hu, and Qiang Ji. Classifier learning with prior probabilities for facial action unit recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5108–5116, 2018.
- [42] Z. Zhang, J. M. Girard, Y. Wu, X. Zhang, P. Liu, U. Ciftci, S. Canavan, M. Reale, A. Horowitz, H. Yang, J. Cohn, Q. Ji, and L. Yin. Multimodal spontaneous emotion corpus for human behavior analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3438–3446, 2016.
- [43] Zheng Zhang, Shuangfei Zhai, and Lijun Yin. Identity-based adversarial training of deep cnns for facial action unit recognition. page 226, 2018.
- [44] K. Zhao, W. Chu, and A. M. Martinez. Learning facial action units from web images with scalable weakly supervised clustering. pages 2090–2099, 2018.
- [45] Y. Zhou, J. Pi, and B. E. Shi. Pose-independent facial action unit intensity regression based on multi-task deep transfer learning. In *Automatic Face & Gesture Recognition (FG 2017), 2017 12th IEEE International Conference on*, pages 872–877, 2017.
- [46] Xinyue Zhu, Yifan Liu, Jiahong Li, Tao Wan, and Zengchang Qin. Emotion classification with data augmentation using generative adversarial networks. *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 349–360, 2018.