Extracting Dynamic Facial Expressions from Naturalistic Videos

Jianxin Wang^{1,2}, Craig Poskanzer^{1,2}, and Stefano Anzellotti¹

¹Boston College

²These authors contributed equally

November 26, 2021

Abstract

Facial expressions are critical in our daily interactions. Studying how humans recognize dynamic facial expressions is an important area of research in social perception, but advancements are hampered by the difficulty of creating well-controlled stimuli. Research on the perception of static faces has made significant progress thanks to techniques that make it possible to generate synthetic face stimuli. However, synthetic dynamic expressions are more difficult to generate; methods that yield realistic dynamics typically rely on the use of infrared markers applied on the face, making it expensive to create datasets that include large numbers of different expressions. In addition, the use of markers might interfere with facial dynamics. In this paper, we contribute a new method to generate large amounts of realistic and well-controlled facial expression videos. We use a deep convolutional neural network with attention and asymmetric loss to extract the dynamics of action units from videos, and demonstrate

that this approach outperforms a baseline model based on convolutional neural networks without attention on the same stimuli. Next, we develop a pipeline to use the action unit dynamics to render realistic synthetic videos. This pipeline makes it possible to generate large scale naturalistic and controllable facial expression datasets to facilitate future research in social cognitive science.

1 Introduction

Facial expressions play a crucial role in our daily interactions and convey information about an individual’s underlying emotional state (Ekman et al., 1987). While contextual information also contributes to emotion inferences (Carroll and Russell, 1996; Hassin et al., 2013; Saxe and Houlihan, 2017; Barrett et al., 2019), facial expressions can help to disambiguate between different possible reactions to the same context, and play an important role in social communication (Schyns et al., 2017). When we give a gift to a friend, we can observe their expression to evaluate whether they are happy or disappointed. When we deliver a talk, we may observe the facial reactions of the audience to infer whether they take issue with the contents. As we navigate a wide range of social situations, we are constantly interpreting facial information with the goal to better understand the attitudes and intentions of those around us. Therefore, studying the mechanisms underlying facial expression recognition is a key step towards understanding how we, as humans, make sense of our social environment.

Some facial expression information can be extracted from a static image of a face (Ekman et al., 1987; Ekman and Heider, 1988; Matsumoto and Ekman, 1989). Static images of facial expressions can be used to disambiguate positive from negative emotions (Ekman and Oster, 1979) (but see Aviezer et al. (2012) for examples of the ambiguity of expressions). In addition, recognition accuracy

using the peak frame from a video (the frame at which the sum of Action Unit activation is highest) is comparable to recognition accuracy when participants see the entire video (Fiorentini and Viviani, 2011). Some studies even reported higher accuracy for expression recognition from still images than from dynamic videos (Widen and Russell, 2015). However, an increasing body of evidence indicates that dynamic information in facial expressions contributes importantly to emotion recognition (Ambadar et al., 2005; Bould et al., 2008; Pitcher and Ungerleider, 2021). Observers can identify subtle properties of expressions from dynamic stimuli, such as the genuineness of a smile (Krumhuber et al., 2013; Dobs et al., 2018). Moreover, videos of facial expressions lead to improved identification of basic emotions in synthetic animations (Hehman et al., 2015). In addition, dynamic expressions are perceived as more intense than static expressions (Biele and Grabowska, 2006), suggesting that dynamic information might affect inferences about the intensity of emotions. Dynamic facial expressions lead to more accurate recognition when the resolution of the images is low (Cunningham and Wallraven, 2009). In real world contexts, this could contribute to expression recognition from a distance. Recent studies suggest that dynamic expressions separate different components of expression mixtures in time (Jack et al., 2014). For example, in the presence of a mixture of surprise and happiness, a dynamic expression could convey information about surprise and about happiness in different time intervals, making it easier to recognize both emotions.

In order to investigate systematically the perception of dynamic expressions, comprehensive and controlled datasets are needed. Naturalistic videos of different individuals displaying the same type of expression can vary not only in the Action Unit (AU) activations, but also in factors such as facial structure, textures, and backgrounds, making it difficult to study the recognition of facial dynamics

controlling for the effect of confounding factors. Developing well-controlled dynamic expression videos is also essential for neuroimaging studies using multi-voxel pattern analysis (MVPA) or representational similarity analysis (RSA, Kriegeskorte et al. (2008)), to investigate the impact of expression dynamics on neural responses controlling for the impact of other factors on the blood-oxygen-level-dependent (BOLD) signal.

Currently, there are at least six different types of dynamic stimuli that can be used to study dynamic face perception: real-world videos of spontaneous expressions, real-world videos of posed expressions, point-light faces, videos from image-based morphing, synthetic facial animations, and photo-realistic face rendering (Dobs et al., 2018). Each of these methodologies comes with a unique set of benefits and limitations.

While real-world videos of spontaneous expressions are the most naturalistic stimuli, datasets with spontaneous expressions are scarce, and often lack some expression categories (e.g. the MMI dataset includes spontaneous expressions but they are limited to disgust, happiness and surprise (Valstar and Pantic, 2010)). Real-world videos of posed expressions typically include a more comprehensive set of emotions, but they offer experimenters limited options to systematically control features like facial structure, facial textures, viewpoint, and lighting. Furthermore, in videos of posed expressions, identity-specific facial dynamics and identity-specific facial structure cannot be manipulated independently (i.e. we generally do not have videos with the same exact AU dynamics performed by different face identities). This makes it difficult to separately study aspects of the dynamics of expressions that are unique to particular individuals independently from the particular facial structure of those individuals.

At the opposite extreme, point-light displays afford precise control and disentangle dynamics from facial structure and textures, but do so by eliminating

elements of realism that might be central to understanding human face perception. Synthetic facial animations, and photo-realistic face rendering in particular, offer a desirable middle ground that affords a combination of realism and control over the generated identity, facial structure, and textures. Synthetic animations can be generated by morphing a face gradually from one expression taken as the starting point, to another expression taken as the ending point. However, morphing-based approaches do not respect the real temporal dynamics of facial expression videos. Dobs et al. (2014) found that given naturalistic and non-naturalistic facial animations, observers prefer the animations containing naturalistic facial dynamics, suggesting that human observers are sensitive to deviations from realistic temporal trajectories of AU activation. Synthetic animations generated by gradual morphing between two expressions do not capture these realistic temporal trajectories.

To address this limitation, recent work has used sensors to measure the activation of facial muscles for a set of actors (Dobs et al., 2018). This approach made it possible to capture more accurate recordings of the dynamics of facial expressions, that were then applied to synthetic faces (Dobs et al., 2018). While this is an important step forward in the generation of controlled dynamic stimuli, the use of reflective markers on the face might affect the facial dynamics of actors for some types of facial expressions. More importantly, this approach is costly and difficult to scale for the generation of large datasets with many identities and expression categories.

In this article, we contribute a technique to create controlled facial expression stimuli with naturalistic dynamics, designed to scale for the production of large datasets. Our method can achieve realism and control by extracting the natural facial dynamics (sequences of the activity of AUs) from videos of real expressions with deep neural networks, and using them to render synthetic

facial expression animations. This approach offers several important advantages for stimulus generation as compared to other techniques such as generative adversarial networks (GANs, Goodfellow et al. (2014)). First, the approach used in this article represents expressions in terms of AUs, facilitating comparison with the prior literature. Second, and more importantly, with GANs it would be difficult to generate videos with the same exact expression dynamics applied to different identities, or with different expression dynamics applied to the same identity. More broadly, the latent representations of GANs representations would not disentangle properties of expression, facial structure, and texture, and as such, lack the compositionality needed to generate datasets with systematic manipulation of cognitively interpretable dimensions.

A recent article introduced a dataset for AU labeling (Yan et al., 2019) and used a CNN to transfer human expressions into virtual avatars. Our method modifies the approach used in (Yan et al., 2019) in two key aspects. First, instead of using a standard CNN, we leverage a network with attentional sampling (Ding et al., 2019). Second, because most AUs are inactive in any given frame, we use a proportional mean square error loss in our model training that makes the net more sensitive to errors affecting *active* action units. This technique prevents the network from being biased towards generating zero outputs for all action units. Thanks to these improvements, our network can label AUs without using a separate face detection software, and outperforms the best architecture in (Yan et al., 2019) on all metrics when tested on the same inputs. The pretrained AU recognition network is available on github: <https://github.com/sccnlab/E-DFE> Finally, we generate the synthetic animations using the publicly available FACSHuman software: <https://github.com/montybot/FACSHuman> (Gilbert et al., 2021). In sum, this article contributes a publicly available, effective and scalable technique for the generation of realistic and controllable dynamic facial

expression videos.

2 Methods

2.1 Overview

We trained a deep neural network model on the FEAFa dataset (Yan et al., 2019). The deep network was able to extract Action Unit (AU) activations from human facial expression videos. The AU activations were used to generate well-controlled facial expression animations using MakeHuman (<http://www.makehumancommunity.org>), that display realistic dynamics.

2.2 FEAFa Dataset

The FEAFa dataset is composed of 123 videos of 122 participants (with one case of two videos for a single participant) recorded in real world conditions. Participants filmed themselves with front-facing cameras performing a series of facial movements and expressions. The average length of a video in the dataset is 32 seconds ($SD = 27$) and the average frame rate is 26 fps ($SD = 6$). Video frames that are not suitable for labelling (due to blurriness etc.) were omitted from our analysis (Yan et al., 2019). In sum, we used a total of 99,356 video frames for our neural network training and testing. Sample video frames are shown in Figure 1 in Yan et al. (2019). In order to track the participants’ facial dynamics, for each video frame, human raters manually scored the values of a set of 24 action units (AUs) defined by the Facial Action Coding System (FACS) (Friesen and Ekman, 1978). Each AU was rated on a continuous scale ranging from 0 (inactive) to 1 (maximally active) (Yan et al., 2019). FACS AUs are parameters that describe facial movements by decomposing them into the motion of specific muscles and face parts (Friesen and Ekman, 1978). These

parameters are widely used in facial expression research, and they are at the basis of software used render facial expression such as FACSHuman (Gilbert et al., 2018). A description of each of the 24 action units rated in the FEAF dataset is reported in the Appendix.

2.3 Deep Neural Network Model

Building a deep neural network to extract the activation of AUs from video frames poses two key challenges. First, information about expressions is localized to specific regions of the face (Gosselin and Schyns, 2001), and the amount of activation for each individual AU is associated with subtle changes. Second, each AU is inactive in most of the images, so there is a risk that deep network models will be biased to always return an activation value of 0. We tackled these challenges with model design choices at the level of the network architecture and of the loss function respectively.

2.3.1 Capturing subtle facial movements

In order to capture the subtle facial movements associated with AUs, we sought a model that could adequately perform this type of fine-grained image recognition. To accomplish this task, we used the Selective Sparse Sampling Network (S3N), which disproportionately samples from highly informative regions of a given image (Ding et al., 2019). The S3N model is built with a pre-trained ResNet50 backbone and trained using our novel proportional mean squared loss (see Loss Function). The purpose of the S3N model is to learn which regions of an input image are highly informative for making classification distinctions between image categories. To accomplish this, the model calculates a class response map M for each AU category, which is calculated by multiplying the values of feature maps extracted from the model’s top convolutional layer by the weight matrix

of a fully-connected layer that uses this feature map to predict AU scores. The peaks of this class response map are thus defined to be extracted features of an image that are highly informative for the prediction task (Ding et al., 2019).

The locations of these peak class response values are collected and partitioned into two sets: discriminative peaks T_d and complimentary peaks T_c , where T_d contains the location of relatively high peak values and T_c has the location of relatively low peak values. T_d and T_c are converted to two sampling maps Q^d and Q^c which are used to selectively sample the input image X such that peak locations of the Q^d and Q^c maps are sampled at a higher rate. These images are re-sampled into discriminative and complimentary images which now contain disproportionately more samples from regions which are highly informative for distinguishing class categories (see Ding et al. (2019) for more details). The two re-sampled images, along with the original input, are taken as input by the ResNet50 backbone to extract the feature representation for each image. The joint feature representation is obtained by the concatenation of feature representations from the original image and the re-sampled images. The final loss is the combination of the loss for each feature representation and the loss for the joint feature representations. For our training, we set the initial learning rate as 0.001 and the weight decay as 0.0001. The S3N model uses image augmentation to better generalize to testing data, however we removed the horizontal flip augmentation given the fact that facial expressions are not necessarily symmetric across the vertical axis.

2.3.2 Addressing the bias for inactive AUs

In a given frame from the FEAFa dataset, the average number of active action units is approximately 3. There are a total of 24 action units per frame, which means that most action units per frame have a value of zero. Because the goal of our study is to map the time courses of active action units, we must use a loss

function that is able to account for this sparse regression problem. A traditional mean squared error loss function treats the loss for each action unit equally, which consequently biases the network to disproportionately predict 0 values for all AUs at each time point. In order to combat this issue, we made use of a ‘proportional mean squared error’ (PMSE) which weights non-zero values (active actions units) more heavily than inactive action units during training. The PMSE function is calculated as:

$$Loss = \sum_{i=1}^n \sum_{j=1}^m \delta_{y_{ij}=0} \lambda_0 (y_{ij} - \hat{y}_{ij})^2 + \delta_{y_{ij} \neq 0} \lambda_a (y_{ij} - \hat{y}_{ij})^2$$

$$\delta_{y_{ij}=0} = \begin{cases} 1, & \text{if } y_{ij} = 0 \\ 0, & \text{otherwise} \end{cases}$$

$$\delta_{y_{ij} \neq 0} = \begin{cases} 1, & \text{if } y_{ij} \neq 0 \\ 0, & \text{otherwise} \end{cases}$$

where n is the number of images per batch and m is the number of action units per image. y_{ij} is the target value for each action unit and $\hat{y}_{ij}$ is the network’s prediction of the corresponding action unit value. Additionally, λ_0 is the weight for inactive units, and λ_a is the weight for active units. They are defined as:

$$\lambda_0 = 1 - \frac{n_0}{t}$$

$$\lambda_a = 1 - \frac{n_a}{t}$$

where n_0 is the number of inactive action units per batch, n_a is the number of active action units per batch, and t is the total number of action units per batch.

2.3.3 Baseline

The performance of the S3N network was compared to the performance of a VGG16 (Simonyan and Zisserman, 2015) architecture pretrained with ImageNET (Deng et al., 2009) and fine-tuned with the FEAFa dataset. This VGG16 architecture was chosen because it achieved the best performance among all models tested in the article that introduced the FEAFa dataset (Yan et al., 2019). We implemented a network architecture identical to the VGG16-based architecture used in (Yan et al., 2019). The VGG16 architecture was fine-tuned using the same FEAFa inputs used to train S3N, and with the same hyperparameters.

2.4 Evaluation metrics

The performance of the deep networks was evaluated using an array of metrics. For each network, we calculated two versions of the mean squared error (MSE). In one version (MSE-im), we divided the sum of square errors by the number of testing images, as a measure of the overall error per image. In the other version (MSE-au), we divided the sum of square errors by the number of actions units, offering a measure of the deviation of the model from the correct labels on an AU by AU basis. In addition, we calculated both MSE-im and MSE-au for active AUs alone, and for inactive AUs alone, to assess the performance of the models on these different sets of AUs separately. This enabled us to evaluate the impact of the difference in base rates between active and inactive AUs.

More specifically, we calculated the mean squared error for all action units over the number of testing images (MSE-im-T) as:

$$\frac{\sum_{i=1}^k \sum_{j=1}^m (y_{ij} - \hat{y}_{ij})^2}{k}$$

where k is the numbers of total testing images.

We also calculated the mean squared error for active action units (MSE-im-A) and non active action units (MSE-im-N) over the number of testing images.

MSE-im-A is:

$$\frac{\sum_{i=1}^k \sum_{j=1}^m \delta_{y_{ij} \neq 0} (y_{ij} - \hat{y}_{ij})^2}{k}$$

and MSE-im-N is:

$$\frac{\sum_{i=1}^k \sum_{j=1}^m \delta_{y_{ij} = 0} (y_{ij} - \hat{y}_{ij})^2}{k}$$

where δ is the Kronecker delta referred before.

2.4.1 MSE-au

To calculate the mean squared error over total number of action units (MSE-au-T), we used the following method:

$$\frac{\sum_{i=1}^k \sum_{j=1}^m (y_{ij} - \hat{y}_{ij})^2}{km}$$

We calculated the mean squared error for active action units (MSE-au-A) and non active action units (MSE-au-N) over total number of action units as well.

MSE-au-A is :

$$\frac{\sum_{i=1}^k \sum_{j=1}^m \delta_{y_{ij} \neq 0} (y_{ij} - \hat{y}_{ij})^2}{\sum_{i=1}^k \sum_{j=1}^m \delta_{y_{ij} \neq 0}}$$

and MSE-au-N is:

$$\frac{\sum_{i=1}^k \sum_{j=1}^m \delta_{y_{ij} = 0} (y_{ij} - \hat{y}_{ij})^2}{\sum_{i=1}^k \sum_{j=1}^m \delta_{y_{ij} = 0}}$$

2.5 Accuracy

In addition to the MSE measures, we calculated a threshold-based accuracy metric. This metric was computed for all AUs (“total AUs accuracy”), and also for the set of active AUs only (“active AUs accuracy”).

Total AUs accuracy on testing data was calculated in the following way:

$$\frac{\sum_{i=1}^k \sum_{j=1}^m \delta_{|y_{ij} - \hat{y}_{ij}| \leq 0.1}}{k}$$

$$\delta_{|y_{ij} - \hat{y}_{ij}| \leq 0.1} = \begin{cases} 1, & \text{if } |y_{ij} - \hat{y}_{ij}| \leq 0.1 \\ 0, & \text{otherwise} \end{cases}$$

A prediction is defined as ‘correct’ if it is within 10% of the target value (a maximum difference of 0.1 in either direction).

Active AUs accuracy in the testing data is calculated in the following way:

$$\frac{\sum_{i=1}^k \sum_{j=1}^m \delta_{y_{ij} \neq 0} \delta_{|y_{ij} - \hat{y}_{ij}| \leq 0.1}}{\sum_{i=1}^k \sum_{j=1}^m \delta_{y_{ij} \neq 0}}$$

2.6 Confusion Matrix

In addition to our measures of accuracy, we also made use of a confusion matrix to better understand the types of errors our model made. In a typical classification task, an input is categorized as belonging to one specific class; if the model incorrectly predicts the input’s class, one can readily see into which class the input was mis-categorized. Thus, a typical confusion matrix denotes the distribution of mis-classified trials for a given class across all other categories, highlighting which classes the model most often confuses. In our task, multiple action units can be active in any given stimuli. As a consequence, an incorrect prediction for one AU does not imply that the model’s other action unit predictions are necessarily all wrong. For example, in a standard classification task if the target value for $AU1 = 1$ then all other AUs must be equal to 0. Thus, if our network predicted $AU1 = 0$ and $AU2 = 1$, then we could correctly conclude that the model predicted a high value for $AU2$ *instead* of $AU1$, because the classes are mutually exclusive. However, in our task, if the target value for $AU1 = 1$, we do not know the value for $AU2$ a priori; thus, if the network predicts $AU1 = 0$ and $AU2 = 1$, then the model could still be correctly predicting the value of $AU2$, even if its prediction for $AU1$ was wrong. To account for this independence across categories, we created a confusion matrix by calculating for a given input, which AU’s were *simultaneously* incorrectly predicted. This means that for each incorrect prediction for a given AU, we can denote the distribution of all other AU’s that are most frequently incorrectly predicted at the same time, This measure provides an indication for the pairs of AUs that the model is most often confusing.

2.7 Rendering synthetic animations with realistic dynamics

We used the MakeHuman software with the FACSHuman plugin (Gilbert et al., 2018) to re-render the dynamic facial expressions. MakeHuman is able to create realistic models of human faces with unique identities. With the help of FACSHuman, we were able to set values for different action units to generate distinct facial expressions in MakeHuman. Figure 1 shows examples of the facial expressions generated in the software.

Figure 1: Facial expressions rendered with MakeHuman

To render naturalistic and controlled facial expression animations, we took the following steps: first, a video of a human participant was divided into its constituent frames. Second, we used the trained model to predict the value of all 24 AUs in each frame. Next, we used a cubic spline with a lambda parameter of 10 in order to smooth the dynamics of each AU over the span of the video. This step was performed because we expect the change in expression to vary gradually over time, and spline interpolation can exploit the gradual change between adjacent timepoints to reduce noisy fluctuation in the AU activation estimates. The resulting interpolated AU time courses were used to create synthetic frames with the MakeHuman software. Finally, we combined the synthetic frames into a

video with the same frame rate as the matching video of the human participant.

One important factor to note was that in the MakeHuman rendering software, left and right are defined with respect to the subject of the image (so the “left” eye is on the right from the perspective of the viewer); however, in the FEAFa dataset, “left” and “right” are defined with respect to the viewer. To counteract this discrepancy, all left/right labels were flipped when rendering the predicted FEAFa expressions in MakeHuman. Unless otherwise specified, throughout the rest of this work, we will use the terms “left” and “right” to refer to the subject of the image’s left and right sides respectively.

Since eyebrows are especially important to distinguish facial expressions (Sadr et al., 2003; Sinha et al., 2006), we took extra care on eyebrow animation. While the FEAFa dataset has individual action units to define the lowering of left vs. right eyebrows, FACS human does not have separate action units to define this left/right eyebrow lowering. To create the effect of lowering the right eyebrow and left eyebrow separately, we generated two frames for each frame in the target video. The first rendered frame lowers both eyebrows using the predicted value of “right eyebrow lowering” AU (Left, with respect to the image subject), and the second frame lowers both eyebrows using the value of the “left eyebrow lowering” AU (Right, with respect to the image subject). Then, we combined the left half of the second frame (subject’s right) with the right half of the first frame (subject’s left) to generate the final version of the frame for the set of AUs.

3 Results

3.1 Mean Square Error

In a first set of analyses, we aimed to evaluate the performance of the network at labeling action units (described in Methods 2.7) in order to get an understanding of the model’s overall ability to capture their activation across the facial topography. To do this, we split our data into a training and testing set, following the established ratio of 80%/20%. In the FEAFa dataset, there are 123 videos. We used 97 videos (78,063 frames) for training and 26 videos (21,293 frames) for testing, yielding a ratio of 79%/21% training and testing data.

We used the model that achieved the best performance in the FEAFa paper as a baseline (a VGG16 architecture pretrained using ImageNET and fine-tuned with the FEAFa dataset) (Yan et al., 2019). We tested our method (a S3N architecture trained with an asymmetric loss) and the baseline method on *non-cropped* video frames, and we computed a set of MSE metrics (see Table 1, and the Methods section for details about how the metrics are computed). MSE calculated using all action units was lower for the S3N model than for the baseline VGG16 model (MSE-im-T = 0.34 for S3N, MSE-im-T = 0.99 for VGG16; MSE-au-T = 0.014 for S3N, MSE-au-T = 0.042 for VGG16). As expected, the S3N model vastly outperformed VGG16 on active action units (MSE-im-A = 0.2 for S3N, MSE-im-A = 0.72 for VGG16; MSE-au-A = 0.081 for S3N, MSE-au-A = 0.289 for VGG16). This might not be too surprising considering that the S3N model was trained using an asymmetric loss that disproportionately penalizes errors on active AUs. However, the S3N model also outperformed the baseline VGG16 model on non-active AUs, despite the fact that errors on these AUs were comparatively less penalized by the asymmetric loss used to train the S3N model (MSE-im-N = 0.14 for S3N, MSE-im-N = 0.27 for VGG16; MSE-au-N = 0.006 for S3N, MSE-au-N = 0.013 for VGG16). In sum, the S3N model trained with

Models	MSE over example			MSE over AUs		
	MSE-im-T	MSE-im-A	MSE-im-N	MSE-au-T	MSE-au-A	MSE-au-N
Ours	0.34	0.2	0.14	0.014	0.081	0.006
Baseline	0.99	0.72	0.27	0.042	0.289	0.013

Table 1: Our method (S3N trained with asymmetric loss) performs better than the baseline (VGG16) in all MSE measures. Overall, predictions for active AUs result in higher MSE values as compared to predictions for non-active AUs prediction. However, our method reduces the gap between active AUs MSE and non-active AUs MSE and improve them significantly comparing to the best model in FEAFa.

the asymmetric loss outperformed the baseline model in all MSE metrics.

3.2 Accuracy

In addition, we further evaluated the performance of our method using a threshold-based accuracy metric (see Methods for details). S3N trained with the asymmetric loss performed better than the baseline VGG16 architecture also in terms of this accuracy metric (see Table 2). It is important to note that the accuracy metric used is stringent. According to this metric, the network’s predictions are deemed to be “accurate” if they are within a 0.1 margin of the target value. As a consequence, in some cases the network correctly identifies that an AU at a given timepoint is active, but the output is still counted as “incorrect” because of a discrepancy in the amount of activation. It is important to note that the inferred AU values could interact with properties of a specific face identity (i.e. the network could infer a greater intensity of the “opening mouth” AU for a person with a larger mouth). For this reason, we additionally evaluated the performance of the S3N network measuring the correlation between the model’s predictions and the changing patterns expressed in the target action units over the span of the video. Rather than focusing on the specific value of the action units at a given time point, the time course data can provide insight into the

net’s ability to extract the dynamics of facial expressions across time (see Figure 2). In fact, when the patterns of target AU activation are plotted alongside the smoothed predictions of our model, it is evident that the fluctuations of AU activation are captured by our network, even though in some cases the model’s prediction are not within 0.1 of the labels produced by human observers (see Figure 2). We further validate our results using the Pearson correlation coefficient of the predicted time course with the ground truth time course. The p-value was calculated using the function `scipy.stats.pearsonr` in the `scipy` library, and it is based on a null hypothesis obtained by sampling correlation values from a β distribution on the interval $[-1, 1]$ with shape parameters $a = b = \frac{n}{2} - 1$, where n is the number of examples. For data with correlation coefficient r , the p-value is the probability that $|r'|$ of a random sample x' and y' drawn from the beta distribution with zero correlation would be greater than or equal to $|r|$.

	Active AUs Accuracy	Total AUs Accuracy
Ours	35%	86%
Baseline	16%	76%

Table 2: This drop in accuracy again reflects the fact that the active action unit regression is more challenging than the non-active action unit regression. Our method outperforms the baseline in both metrics.

Time course data for 2 representative action units (“jaw drop” and “right eye closure”) in two different randomly selected FEAFa participants (subject number 20 and 40) are shown in Figure 2. The dynamics of each AU are tracked by the network’s predictions. In fact the network’s predictions showed a strong correlation with the target values for each action individual AU (Jaw Drop Subject 20: $R = 0.99$, $P < 0.0001$; Jaw Drop Subject 40: $R = 0.81$, $P < 0.0001$; Subject 20 Right Eye Closure $R = 0.96$, $P < 0.0001$; Subject 40 Right Eye Closure $R = 0.84$, $P < 0.0001$).

Figure 2: Changes in activation over time in the target action units are captured by our model, despite differences in the overall magnitude of the activation. Given that the network has an increased accuracy for non-active action units, we can observe, across all videos, that the predicted values tightly track the target AU values in spans where the unit has a target value of 0. For active AUs, while the measured accuracy is lower, the network is still able to model the patterns of activation over time, with high correlations between the predicted AU intensities and the ground truth.

3.3 Structure of the Errors

While errors are inevitable in any classification task, the *types* of mistakes that our model makes can be especially informative of its ability to correctly identify and accurately replicate fine-grained facial movements. A confusion matrix is a commonly used tool in classification tasks in which one can examine whether the model mistakenly switches the assignments for empirically related or unrelated

categories. We apply a confusion matrix to check what types of AUs our model is confusing and whether the confused AUs are “related”. In order to define “related” AUs, we note that in MakeHuman, AUs are grouped in subsets defined by their proximity and relationship to each other; in this way, we categorize “related” AUs, as AUs belonging to the same MakeHuman delineated classification (see Appendix, MakeHuman AU Categories). In order to probe our model’s incorrect classifications, we made use of a novel confusion matrix which allows us to study simultaneously mis-classified AUs as a measure of model confusion (see Methods; Confusion Matrix). In our analysis, we found several pairs of AUs that had a greater likelihood of being simultaneously incorrect; specifically. The confusion matrix illustrates that there is a high co-occurrence of mistakes between the left/right AUs for most symmetric features, most evident in the eye close action units. Moreover, our model also often “confuses” the left/right brow lower with the left/right Eye close action units as well as mistaking the left/right brow raise with the left/right lid raise. Empirically, these sets of mistakes (which co-occur approximately 1000 times more often than any other sets of errors) reveal an underlying relationship between AUs in close proximity as compared to AUs governing completely separate topographical sections of the face. In fact, the AUs that our network most commonly confuses all belong to the same FACS subset of AUs (“Upper Face AUs”).

Figure 3: Using a novel confusion matrix developed for continuous multidimensional outputs (see Methods, Confusion Matrix) we can observe the co-occurrence of incorrect predictions for pairs of AUs. Left and right eye close AUs were the most confused AUs.

3.4 Rendered Expression

Due to privacy concerns, we did not use the videos from the FEFA dataset as an example to illustrate our rendered animation. Instead, we created our own facial expression sample video to display the rendering results. The sample video is 10 second long and consisted of four different randomly selected facial expressions: no facial movement, pout, right lip corner raise with closed right eye, and left lip corner raise with closed left eye. After feeding the video frames into our pretrained model, we generated a rendered facial animation in MakeHuman from network’s predictions. As shown in Figure 4, four selected rendered animation

frames have facial movements that match the corresponding sample video frames.

Figure 4: Animations rendered with MakeHuman using the AU intensities extracted with our S3N network (bottom row) match the expressions in the original video (top row).

4 Discussion

In this article we introduced an approach to render well-controlled animations of facial expressions with naturalistic facial dynamics. First, we trained a deep neural network using sparse-selective sampling (S3N) and an asymmetric loss developed specifically for the extraction of action unit activations. Next, we demonstrated that this network outperforms other alternatives such as VGG16. Finally, we used the network’s outputs in conjunction with FACSHuman to render realistic facial expression animations.

The pre-trained network is freely available on this github page: <https://github.com/sccnlab/E-DFE>. It can be used by researchers studying the cognitive mechanisms for face perception (and their underlying neural bases) to generate stimuli that are well-controlled in terms of lighting and viewpoint, as well as other low-level properties that are more difficult to control with video editing (e.g. textures). Importantly, this approach enables researchers to

independently manipulate face structure and expression dynamics; it is possible to apply the same expression dynamics to different face identities, making it easier to investigate how changes in facial structure and changes in expression dynamics separately affect recognition accuracy. Unlike previous techniques developed to generate controlled dynamic expression stimuli (Dobs et al., 2018), the approach introduced in this article does not require hiring actors to perform the expressions and measuring facial dynamics with reflective markers. In this way, the present approach is more economical and easier to scale to large numbers of videos.

The method we introduced outperformed the baseline VGG16 architecture on all metrics used. The improvement was greater for extracting the intensity of active action units, likely thanks to the use of an asymmetric loss that disproportionately penalized errors on active AUs. This loss was chosen because at each given time most AUs are inactive, leading to the risk of bias in favor of inactive AUs (due to the different base rates). Importantly, despite the use of an asymmetric loss, the S3N network still outperformed VGG16 also when measuring the MSE on inactive AUs alone. This shows that the S3N is able to accurately identify the AUs that are inactive.

As the accuracy of our model achieves a score of 85% at labeling the action units, our model is performing well in capturing general dynamic of facial expressions. Although our model is able to accurately map the changes in AU activity, there are still several video frames for which the model’s prediction are not within 0.1 of the ground truth. One factor that could contribute to the gaps between our predictions and the target values is the natural variations in facial structures between individuals. For example, the size of a given subject’s mouth can affect the predicted activity for mouth related AUs. In one confusing case, a small opening of a larger mouth might appear equivalent to a large opening of a

smaller mouth. In this scenario, our model may have trouble disentangling the correct AU values of the given face, while still accurately mapping the changes in AU activity as the facial features move over time. In line with this possibility, correlations between predicted AU timecourses and the ground-truth timecourse are high, with r-values above 0.8 for the examples tested.

While our method outperforms previous approaches, the present method can be improved in several respects. First, cropping the video frames to the region of the face before entering them as inputs to the network could further improve performance. However, using non-cropped images is less computationally intensive and removes one step in the generation of the stimuli, thus making this approach more user friendly. This can help to encourage a broader adoption of this method. Second, the present method is based on extracting AU values from individual frames, and only uses the temporal smoothness of facial dynamics to interpolate the time series of outputs. More accurate results might be achieved using neural network architectures that directly leverage the temporal dependency between subsequent frames (i.e. see Wei et al. (2021)), such as LSTM. These approaches might be especially effective at predicting the AU activations for blurry frames.

Taking this work as a starting point, there are several promising future directions to improve the current technology. First, large numbers of rendered videos with naturalistic dynamics could be used as a form of data augmentation, to further improve the accuracy of AU predictions. Second, the AU timecourses extracted from large datasets of naturalistic videos could be used to train generative adversarial networks (GANs) that would then be able to produce novel but realistic expression timecourses. Already as it is, the present method paves the way to new lines of research in face perception. Thanks to the efficient extraction of AU timecourses, it is now possible to investigate the structure of

the space of dynamic expressions, extending the work that has been done to identify a space for face identity (Valentine et al., 2016; Carlin and Kriegeskorte, 2017). In addition, because naturalistic AU timecourses can be systematically modified, it is possible to study the representation of expression dynamics by testing psychophysical detection thresholds for different kinds of perturbations and how they interact with differences in face identity. In sum, this work offers a variety of options to create and manipulate dynamic expression animations, that can help address new questions about the cognitive mechanisms underlying face perception.

5 Appendix

AU	AU name	AU name in FACS	AU	AU name	AU name in FACS
1	Left Eye Close	Upper Face AUs/L43	13	Right Lip Corner Pull	Lower Face AUs/R12
2	Right Eye Close	Upper Face AUs/R43	14	Left Lip Corner Stretch	Lower Face AUs/L20
3	Left Lid Raise	Upper Face AUs/L5	15	Right Lip Corner Stretch	Lower Face AUs/R20
4	Right Lid Raise	Upper Face AUs/R5	16	Upper Lip Suck	Lower Face AUs/28_a
5	Left Brow Lower	Upper Face AUs/4	17	Lower Lip Suck	Lower Face AUs/28_bottom
6	Right Brow Lower	Upper Face AUs/4	18	Jaw Thrust	Miscellaneous AUs/29
7	Left Brow Raise	Upper Face AUs/L2	19	Upper Lip Raise	Lower Face AUs/10
8	Right Brow Raise	Upper Face AUs/R2	20	Lower Lip Depress	Lower Face AUs/16
9	Jaw Drop	Lip Parting and Jaw Opening/26	21	Chin Raise	Lower Face AUs/17
10	Jaw Slide Left	Miscellaneous AUs/L30	22	Lip Pucker	Lower Face AUs/18
11	Jaw Slide Right	Miscellaneous AUs/R30	23	Cheeks Puff	Miscellaneous AUs/34
12	Left Lip Corner Pull	Lower Face AUs/L12	24	Nose Wrinkle	Lower Face AUs/9

Figure 5: Descriptions of the Action Units that are labeled in the FEAFa dataset.

Upper Face AUs	Lower Face AUs	Lip Parting and Jaw Opening	Head Positions	Miscellaneous AUs
Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%
1 Inner Brow Raise	9 Nose Wrinkle	25 Lips Part	51 Head Turn Left	19 Tongue Show
Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%
1 Inner Brow Raise (Left)	9 Nose Wrinkle (Alternative A)	25 Lips Part (Left)	52 Head Turn Right	29 Jaw Thrust
Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%
1 Inner Brow Raise (Right)	9 Nose Wrinkle (Left)	25 Lips Part (Right)	53 Head Up	30 Jaw Sideways (Left)
Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%
2 Outer Brow Raise	9 Nose Wrinkle (Right)	25 Lips Part (Alternative A)	54 Head Down	30 Jaw Sideways (Right)
Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%
2 Outer Brow Raise (Alternative)	10 Upper Lip Raiser	26 Jaw Drop	55 Head Tilt Left	Intensity: 0%
Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%
2 Outer Brow Raise (Left)	10 Upper Lip Raiser (Left)	26 Jaw Drop Bottom Lip Down (Alternative)	56 Head Tilt Right	31 Jaw Clencher
Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%
2 Outer Brow Raise (Right)	10 Upper Lip Raiser (Right)	26 Jaw Drop Tongue Down (Alternative)	57 Head Forward	Intensity: 0%
Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%	32 Bile
4 Brow Lowerer	11 Nasolabial Furrow Deepener	26 Jaw Drop Tongue Out (Alternative)	58 Head Back	Intensity: 0%
Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%	33 Blow
4 Brow Lowerer (Alternative A)	11 Nasolabial Furrow Deepener (Left)	26 Jaw Drop (Alternative)	Intensity: 0%	Intensity: 0%
Intensity: 0%	Intensity: 0%	17 Chin Raiser (Alternative A)	Intensity: 0%	34 Puff
4 Brow Lowerer (Alternative B)	11 Nasolabial Furrow Deepener (Right)	Intensity: 0%	Intensity: 0%	Intensity: 0%
Intensity: 0%	Intensity: 0%	Intensity: 0%	Intensity: 0%	35 Cheek Suck
5 Upper Lip Raise	12 Lip Corner Puller	18 Lip Pucker	Intensity: 0.60%	Intensity: 0%
Intensity: 0%	Intensity: 0%	Intensity: 0%	61 Eye Left (Right eye)	36 Nostril Dilate
5 Upper Lip Raise (Left)	12 Lip Corner Puller (Alternative A)	18 Lip Pucker (Alternative)	Intensity: 0%	Intensity: 0%
Intensity: 0%	Intensity: 0%	20 Lip Stretch	61 Eye Left (Left eye)	39 Nostril Compress
5 Upper Lip Raise (Right)	12 Lip Corner Puller (Left)	Intensity: 0%	Intensity: 0%	Intensity: 0%
Intensity: 0%	Intensity: 0%	Intensity: 0%	62 Eyes Right	Both eyes pupils contract
6 Cheek Raise	12 Lip Corner Puller (Right)	20 Lip Stretch (Left) (Alternative)	62 Eyes Right (Left eye)	Intensity: 0%
Intensity: 0%	Intensity: 0%	Intensity: 0%	62 Eyes Right (Right eye)	Both eyes pupils dilate
6 Cheek Raise (Alternative A)	13 Sharp Lip Pucker	20 Lip Stretch (Right) (Alternative)	Intensity: 0%	Intensity: 0%
Intensity: 0%	Intensity: 0%	Intensity: 0%	63 Eye Up	Left eye pupil contract
6 Cheek Raise (Alternative B)	13 Sharp Lip Pucker (Left)	22 Lip Funneler (Bottom Lip)	63 Eye Up (Left eye)	Intensity: 0%
Intensity: 0%	Intensity: 0%	Intensity: 0%	63 Eye Up (Right eye)	Left eye pupil dilate
6 Cheek Raise (Left)	13 Sharp Lip Pucker (Right)	22 Lip Funneler (Both Lip)	64 Eyes Down	Right eye pupil contract
Intensity: 0%	Intensity: 0%	Intensity: 0%	64 Eye Down (Left eye)	Intensity: 0%
6 Cheek Raise (Right)	14 Dimpler	23 Lip Tightener	64 Eye Down (Right eye)	Right eye pupil dilate
Intensity: 0%	Intensity: 0%	Intensity: 0%	65 Wall eyes	Intensity: 0%
7 Lid Tight	14 Dimpler (Left)	24 Lip Pressor	66 Cross eyes	
Intensity: 0%	Intensity: 0%	Intensity: 0%		
7 Lid Tight (Left)	15 Lip Corner Depressor	28 Lips Suck		
Intensity: 0%	Intensity: 0%	Intensity: 0%		
7 Lid Tight (Right)	15 Lip Corner Depressor (Left)	28 Lips Suck (Upper lip)		
Intensity: 0%	Intensity: 0%	Intensity: 0%		
43 Eye Closure	15 Lip Corner Depressor (Right)	28 Lips Suck (Lower lip)		
Intensity: 0%	Intensity: 0%			
43 for 46 Blink, 46 Wink (Left)	16 Lower Lip Depress			
Intensity: 0%	Intensity: 0%			
43 for 46 Blink, 46 Wink (Right)	17 Chin Raiser			
Intensity: 0%				

Figure 6: MakeHuman AU Categories. There are six different AU categories in the table: Upper Face AUs, Lower Face AUs, Lip Parting and Jaw Opening, Head Positions, Eye Positions, and Miscellaneous AUs

Acknowledgments

This work was supported by NSF CAREER award #1943862 to S.A. We would also like to thank Yan et al. (2019) for making the FEAF dataset available for our use.

References

- Ambadar, Z., Schooler, J. W., and Cohn, J. F. (2005). Deciphering the enigmatic face: The importance of facial dynamics in interpreting subtle facial expressions. *Psychological science*, 16(5):403–410.
- Aviezer, H., Trope, Y., and Todorov, A. (2012). Body cues, not facial expressions, discriminate between intense positive and negative emotions. *Science*, 338(6111):1225–1229.
- Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M., and Pollak, S. D. (2019). Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological science in the public interest*, 20(1):1–68.
- Biele, C. and Grabowska, A. (2006). Sex differences in perception of emotion intensity in dynamic and static facial expressions. *Experimental brain research*, 171(1):1–6.
- Bould, E., Morris, N., and Wink, B. (2008). Recognising subtle emotional expressions: The role of facial movements. *Cognition and Emotion*, 22(8):1569–1587.
- Carlin, J. D. and Kriegeskorte, N. (2017). Adjudicating between face-coding models with individual-face fmri responses. *PLoS computational biology*, 13(7):e1005604.
- Carroll, J. M. and Russell, J. A. (1996). Do facial expressions signal specific emotions? judging emotion from the face in context. *Journal of personality and social psychology*, 70(2):205.
- Cunningham, D. W. and Wallraven, C. (2009). Dynamic information for the recognition of conversational expressions. *Journal of Vision*, 9(13):7–7.

- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee.
- Ding, Y., Zhou, Y., Zhu, Y., Ye, Q., and Jiao, J. (2019). Selective sparse sampling for fine-grained image recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6599–6608.
- Dobs, K., Bülthoff, I., and Schultz, J. (2018). Use and usefulness of dynamic face stimuli for face perception studies—a review of behavioral findings and methodology. *Frontiers in psychology*, 9:1355.
- Dobs, K., Bülthoff, I., Breidt, M., Vuong, Q. C., Curio, C., and Schultz, J. (2014). Quantifying human sensitivity to spatio-temporal information in dynamic faces. *Vision Research*, 100:78—87.
- Ekman, P., Friesen, W. V., O’Sullivan, M., Chan, A., Diacoyanni-Tarlatzis, I., Heider, K., Krause, R., LeCompte, W. A., Pitcairn, T., Ricci-Bitti, P. E., Scherer, K., Tomita, M., and Tzavaras, A. (1987). Universals and cultural differences in the judgments of facial expressions of emotion. *Journal of Personality and Social Psychology*, 53(4):712–717.
- Ekman, P. and Heider, K. G. (1988). The universality of a contempt expression: A replication. *Motivation and Emotion*, 12:303–308.
- Ekman, P. and Oster, H. (1979). Facial expressions of emotion. *Annual review of psychology*, 30(1):527–554.
- Fiorentini, C. and Viviani, P. (2011). Is there a dynamic advantage for facial expressions? *Journal of Vision*, 11(3):17–17.
- Friesen, E. and Ekman, P. (1978). Facial action coding system: a technique for the measurement of facial movement. *Palo Alto*, 3(2):5.

- Gilbert, M., Demarchi, S., and Urdapilleta, I. (2018). Facshuman a software to create experimental material by modeling 3d facial expression. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents*, pages 333–334.
- Gilbert, M., Demarchi, S., and Urdapilleta, I. (2021). Facshuman, a software program for creating experimental material by modeling 3d facial expressions. *Behavior Research Methods*, pages 1–21.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial networks. *arXiv preprint arXiv:1406.2661*.
- Gosselin, F. and Schyns, P. G. (2001). Bubbles: a technique to reveal the use of information in recognition tasks. *Vision research*, 41(17):2261–2271.
- Hassin, R. R., Aviezer, H., and Bentin, S. (2013). Inherently ambiguous: Facial expressions of emotions, in context. *Emotion Review*, 5(1):60–65.
- Hehman, E., Flake, J. K., and Freeman, J. B. (2015). Static and dynamic facial cues differentially affect the consistency of social evaluations. *Personality and Social Psychology Bulletin*, 41(8):1123–1134.
- Jack, R. E., Garrod, O. G., and Schyns, P. G. (2014). Dynamic facial expressions of emotion transmit an evolving hierarchy of signals over time. *Current biology*, 24(2):187–192.
- Kriegeskorte, N., Mur, M., and Bandettini, P. A. (2008). Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2:4.
- Krumhuber, E. G., Kappas, A., and Manstead, A. S. (2013). Effects of dynamic aspects of facial expressions: A review. *Emotion Review*, 5(1):41–46.

- Matsumoto, D. and Ekman, P. (1989). American-japanese cultural differences in intensity ratings of facial expressions of emotion. *Motivation and Emotion*, 13:143–157.
- Pitcher, D. and Ungerleider, L. G. (2021). Evidence for a third visual pathway specialized for social perception. *Trends in Cognitive Sciences*, 25(2):100–110.
- Sadr, J., Jarudi, I., and Sinha, P. (2003). The role of eyebrows in face recognition. *Perception*, 32(3):285–293.
- Saxe, R. and Houlihan, S. D. (2017). Formalizing emotion concepts within a bayesian model of theory of mind. *Current opinion in Psychology*, 17:15–21.
- Schyns, P. G. et al. (2017). Toward a social psychophysics of face communication.
- Simonyan, K. and Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*.
- Sinha, P., Balas, B., Ostrovsky, Y., and Russell, R. (2006). Face recognition by humans: Nineteen results all computer vision researchers should know about. *Proceedings of the IEEE*, 94(11):1948–1962.
- Valentine, T., Lewis, M. B., and Hills, P. J. (2016). Face-space: A unifying concept in face recognition research. *The Quarterly Journal of Experimental Psychology*, 69(10):1996–2019.
- Valstar, M. and Pantic, M. (2010). Induced disgust, happiness and surprise: an addition to the mmi facial expression database. In *Proc. 3rd Intern. Workshop on EMOTION (satellite of LREC): Corpora for Research on Emotion and Affect*, page 65. Paris, France.

- Wei, C., Lu, K., Gan, W., and Xue, J. (2021). Spatiotemporal features and local relationship learning for facial action unit intensity regression. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 1109–1113.
- Widen, S. C. and Russell, J. A. (2015). Do dynamic facial expressions convey emotions to children better than do static ones? *Journal of Cognition and Development*, 16(5):802–811.
- Yan, Y., Lu, K., Xue, J., Gao, P., and Lyu, J. (2019). Feafa: A well-annotated dataset for facial expression analysis and 3d facial animation. In *2019 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, pages 96–101. IEEE.