

Semi-supervised Facial Action Unit Intensity Estimation with Contrastive Learning

Enrique Sanchez, Adrian Bulat, Anestis Zaganidis, and Georgios Tzimiropoulos

Samsung AI Center, Cambridge, UK

{e.lozano, adrian.bulat, a.zaganidis, georgios.t}@samsung.com

Abstract. This paper tackles the challenging problem of estimating the intensity of Facial Action Units with few labeled images. Contrary to previous works, our method does not require to manually select key frames, and produces state-of-the-art results with as little as 2% of annotated frames, which are *randomly chosen*. To this end, we propose a semi-supervised learning approach where a spatio-temporal model combining a feature extractor and a temporal module are learned in two stages. The first stage uses datasets of unlabeled videos to learn a strong spatio-temporal representation of facial behavior dynamics based on contrastive learning. To our knowledge we are the first to build upon this framework for modeling facial behavior in an unsupervised manner. The second stage uses another dataset of randomly chosen labeled frames to train a regressor on top of our spatio-temporal model for estimating the AU intensity. We show that although backpropagation through time is applied only with respect to the output of the network for extremely sparse and randomly chosen labeled frames, our model can be effectively trained to estimate AU intensity accurately, thanks to the unsupervised pre-training of the first stage. We experimentally validate that our method outperforms existing methods when working with as little as 2% of randomly chosen data for both DISFA and BP4D datasets, without a careful choice of labeled frames, a time-consuming task still required in previous approaches.

Keywords: Semi-supervised learning, Unsupervised representation learning, Facial Action Units

1 Introduction

Facial actions are one of the most important means of non-verbal communication, and thus their automatic analysis plays a crucial role in making machines understand human behavior. The set of facial actions and their role in conveying emotions has been defined by the Facial Action Coding System (FACS [7]). FACS define a set of atomic facial movements, known as Action Units, whose combination can be correlated with both basic and complex emotions. Action Units are categorically modeled according to their intensities, with values that range from 0, indicating the absence of an AU, to 5, indicating the maximum level of expressivity of an AU.

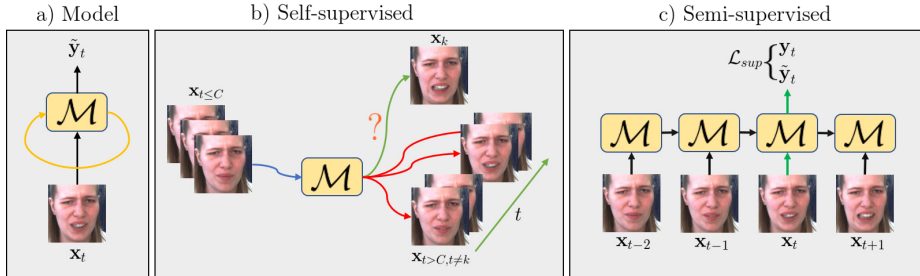


Fig. 1. Overview of our proposed approach. a) We want to learn a temporal model for facial Action Unit intensity estimation that can be learned with a few *randomly chosen* set of annotations. b) We propose to first learn a strong feature representation through *contrastive learning*, by training the model to solve a proxy task, that requires no labeling. The proxy task uses a simple non-linear predictive function, and aims at predicting the feature representation that is similar to that of a specific target frame \mathbf{x}_k , and at the same time dissimilar to the rest of the given sequence. c) Once the model is trained to produce rich facial feature representations, we can further train it for facial Action Unit intensity estimation by using sequences in which the labeled frames are in random position in the sequence.

While plenty of works exist for AU detection (i.e. whether an AU occurs or not in an image, regardless of its intensity), the more challenging task of automatically estimating their intensity has received less attention. Recent advances in supervised methods incorporate a variety of techniques including attention [41, 51, 64], co-occurrence modeling [47, 59], or temporal dynamics [20, 4]. All these methods however require a large number of training instances to work properly, which entails the time consuming task of dataset labeling. This is even more profound for the problem of AU detection and intensity estimation where labeling is typically performed at a frame-level (i.e. each video frame must be labeled). Our goal in this paper is to devise a method that can effectively estimate AU intensity even when a very small (of the order of 2%) and randomly chosen set of frames is used for model training.

This line of work has been pursued by the research community only recently [30, 59, 63, 60, 62, 61]. While these works have shown remarkable results, they still have some limitations: 1) They work on a per-frame basis, by learning a strong image-based feature representation that can later be used for AU intensity estimation. While some of these methods impose a temporal smoothing through ranking [62, 60], they still aim at learning a per-frame representation. 2) Some of these methods also require having a very specific set of annotations in hand [63, 60, 62, 61]. In particular, they work by assuming that annotations are available for “peak” and “valley” frames, i.e. frames corresponding to a local maximum or minimum on the intensity. Identifying these frames requires a qualitative labeling of sequences before the annotation step. That is, while using peak and valley frames is effective with as little as 2% of annotated frames, these anno-

tations require indeed evaluating segments of videos which is also an expensive operation.

In this paper, we take a different path in semi-supervised AU intensity estimation. We firstly assume a practical setting where only a very sparse (of the order of 2%) and randomly chosen set of frames (of a given dataset) is labeled. Our aim is to train a model which, even if trained on this sparse set of annotated frames, it can make dense per-frame predictions of AU intensity for all frames in a given test sequence. To this end, we build a model that combines a feature extraction network (ResNet-18, [17]) with a GRU unit [3], and a regressor head on top of the GRU that can make predictions for each frame and train it using back-propagation through time using only the predictions of the network at the sparsely labeled frames. We found however that training this model from scratch in an end-to-end manner using only a small number of labeled frames is a rather difficult task. Hence, we further propose to firstly train the backbone (i.e. the feature extractor and the GRU unit) in an unsupervised way on different unlabeled datasets, using the recent framework of contrastive learning [36, 54, 52, 15, 18, 43, 1]. The backbone is trained end-to-end with a contrastive loss on unlabeled videos of facial behavior, and used thereafter to train a model for AU intensity estimation using only few labels. An overview of our approach is shown in Fig. 1.

Our **main contributions** can be summarized as follows:

- We are the first to propose a temporal modeling of facial actions that can be learned with a sparse set of discontinuously and randomly chosen annotated facial images. Our practical approach to AU intensity estimation requires as little as 2% of annotated data.
- We are the first to apply the framework of contrastive learning for semi-supervised AU intensity estimation. We propose a two-stage pipeline where a model for obtaining a spatio-temporal facial representation is firstly trained on large unlabeled datasets of facial behavior using a contrastive learning formulation, and, then, the model is effectively trained for the task of AU intensity estimation using a small number of sparsely annotated video frames.
- Our approach achieves state-of-the-art results on both BP4D and DISFA, when using a randomly chosen subset of 2% frames.

2 Related work

2.1 Action Unit modeling

The majority of existing works in facial Action Unit intensity estimation work on a fully supervised way, i.e. by assuming that a large amount of labeled data is available [31, 9]. Existing supervised methods are often split into methods that exploit the geometric structure of faces, also referred to as patch-based [64, 29, 65], methods that exploit the temporal correlation of AUs [20, 4], and those that exploit the correlation that exists between different Action Units [47, 10, 39, 38, 40]. Other methods attempt to exploit different types of correlation [33, 28, 5, 53,

49, 8, 51]. Most recent methods reporting state of the art results on Action Unit intensity estimation build on AutoEncoders [44], or on Heatmap Regression [41, 35].

While there is a vast amount of literature in fully supervised methods for Action Unit intensity estimation, few works have focused on the more challenging task of addressing the same goal in a semi-supervised manner. In [30], a twin autoencoder is used to disentangle facial movements from head pose, the learned representation is then used for Action Unit detection. In [63], an ordinal relevance regression method is applied. In [59], the relation between emotions and Action Units is used to generate a knowledge-graph that allows the use of the emotion labels to train the AU detector without labels. In [48] a Restricted Boltzmann Machine and Support Vector Regression approach to model the AUs is proposed. In [60], a knowledge-based approach is proposed, exploiting the temporal variation that exists between peak and valley frames. In [62], the temporal ranking is exploited in a similar way, and a novel ADMM approach is used to enforce different constraints. In [61], a learnable context matrix is used for each AU, which combined with several patch-based feature fusion is capable of learning AU intensity from key frames only. The majority of these methods however impose temporal constraints based on the fact that annotations are available for *peak* and *valley* frames, i.e. frames that correspond to a local maximum and minimum, respectively. Our method bypasses this need by applying a self-supervised pre-learning step, that learns a strong feature representation. Our approach uses as little as 2% of annotated data. However, contrary to these methods, we use a *randomly chosen* set of data.

2.2 Self-supervised learning in Computer Vision

Self-supervised learning, often referred to as unsupervised learning, is an active research topic in Machine Learning. It involves defining a pretext or proxy task, with a corresponding loss, that leads to strong feature representation. This pretext task does not require the data to be labeled, e.g. it can be predicting relative location of image patches [6], predicting rotation from images [11], sorting frames in a shuffled video [34], denoising [46] or colorization [56], or pseudo-labeling through clustering [2].

More recently, a number of state-of-the-art self-supervised methods based on the so-called contrastive learning formulation have been proposed [36, 52, 54, 15, 18, 43, 14, 1]. The idea is to define a loss for unsupervised learning which maximizes the similarity between the feature representations of two different instances of the same training sample while simultaneously minimizes the similarity with the representations computed from different samples. Among these methods, in this work, we build upon the Contrastive Predictive Coding of [36] which allows for a temporal model to be learned in an unsupervised manner from video data, and hence it is particular suitable for modeling facial behavior dynamics. To our knowledge, we are the first building upon this model for unsupervised modeling of facial behavior. We also show how to apply this model for semi-supervised AU intensity estimation.

3 Method

We are interested in learning a spatio-temporal model for facial Action Unit intensity estimation, that is capable of modeling the temporal dynamics of expressions, as well as their feature representation, and that can be learned with few labels. Thus, we first introduce the problem statement and notation in Sec. 3.1, then we devote Sec. 3.2 to describing the model that will produce a structured output representation, and present our two-stage approach to the model learning in Sec. 3.3 and Sec. 3.4.

3.1 Problem statement and notation

We are interested in learning a model capable of predicting the intensity of some Action Units in a given sequence, that at the same time captures the spatial features responsible of displayed expressions and models the temporal correlation between them. We want such a model to be learned in a scenario where only a small set of frames are annotated with Action Unit intensity.

Let $\mathbf{X} = \{\mathbf{x}_t\}_{t=1}^T$ be a sequence of T video frames, where $\mathbf{x}_t \in \mathbb{R}^{3 \times H \times W}$ is an RGB image of size $H \times W$. Our goal is to learn a model \mathcal{M} that produces a structured output $\tilde{\mathbf{Y}} = \{\tilde{\mathbf{y}}_t\}_{t=1}^T$, where each $\tilde{\mathbf{y}}_t \in \mathbb{R}^N$ corresponds to the predicted intensity of each of the N Action Units of interest. In a fully supervised setting, we would be given a set of ground-truth labels \mathbf{y}_t for each sequence, and thus the model $\mathcal{M}(\cdot; \theta)$ could be learned through regression and backpropagation through time [50]. However, learning the model \mathcal{M} is a hard task, and requires a vast amount of per-frame labeled sequences. Also, training deep temporal models is often a challenging task, very sensitive to the initialization. We now describe our approach to learning the model efficiently with few labels.

3.2 Network

Our model \mathcal{M} is split into a feature extraction block $f(\cdot; \theta_f)$, a temporal module $g(\cdot; \theta_g)$, and a final regressor head $c(\cdot; \theta_c)$, with θ the parameters of each module. The choice for the *feature extractor* $f(\cdot; \theta_f)$ is that of a ResNet-18 [17], without the last fully connected layer. In order to keep a locally-based receptive field for the extracted features, we remove the average pooling layer. To reduce the complexity of the network, the last convolutional block is set to 256 channels rather than the typical 512. This way, for an input resolution of 128×128 , the network produces a feature representation of 256 channels, and of a spatial resolution of 8×8 , i.e. $f(\cdot; \theta_f) \in \mathbb{R}^{256 \times 8 \times 8}$. This way, each spatial output *attends* a region of 16×16 pixels in the input image.

The form of the *temporal block* $g(\cdot; \theta_g)$ is that of a Convolutional Gated Recurrent Unit (GRU [3]), with the latent dimension set to be the same as for the feature representation, i.e. $256 \times 8 \times 8$. The GRU receives, at time t , the input feature representation $f(\mathbf{x}_t; \theta_f) \in \mathbb{R}^{256 \times 8 \times 8}$, along with the previous hidden state $\mathbf{h}_{t-1} \in \mathbb{R}^{256 \times 8 \times 8}$, and propagates the hidden state at time t as $\mathbf{h}_t = g(f(\mathbf{x}_t; \theta_f), \mathbf{h}_{t-1}; \theta_g) \in \mathbb{R}^{256 \times 8 \times 8}$. The GRU is modified to receive a tensor

rather than a vector, by setting all the internal gates to be convolutional layers, with kernel size 1×1 .

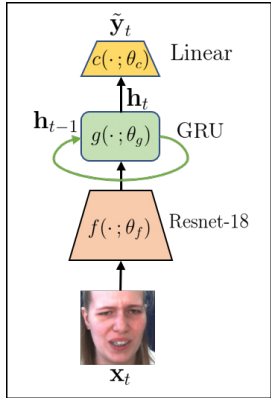


Fig. 2. Network: the model consists of a *feature extractor*, a *temporal block*, and a *regressor head*.

Our learning goal is defined through a proxy *predictive function* $p_k(\cdot; \theta_p)$, tasked with predicting the future feature representations $f(\mathbf{x}_{t+k}; \theta_f)$, at some time $t+k$ from a given contextual information $g(f(\mathbf{x}_t); \theta_g)$, computed up to time t . The learning goal is to make $p_k(g(f(\mathbf{x}_t)))^1$ similar to $f(\mathbf{x}_{t+k})$, and at the same time different from other feature representations computed at the same time step k for a different image \mathbf{x}' , $f(\mathbf{x}'_{t+k})$, at the same image \mathbf{x} but at a different time step k' , $f(\mathbf{x}_{t+k'})$, and at a different image \mathbf{x}' and time step k' , $f(\mathbf{x}'_{t+k'})$. With p being a simple non-linear function, the learning burden lies on a feature representation capable of predicting the future. While in [36] the time step k is fixed, in [14] the time $t+k$ is added recursively from $t+1$ to $t+k$.

Put formally, for a sequence $\mathbf{X} = \{\mathbf{x}_t\}_{t=1}^T$, the first C frames will represent the *context*, whereas the last $P = T - C$ frames will be used for the *predictive* task. Throughout our experiments, $T = 15$ frames, and C and P are set to 10 and 5 frames, respectively, i.e. from the context estimated for the first 10 frames, the goal is to predict the next 5 feature representations.

Let $\mathbf{f}_t^{(i,j)}$ represent the feature representation $f(\mathbf{x}_t)$ at time t , at the spatial location (i, j) . Similarly, let $\mathbf{p}_{t+k}^{(i,j)}$ represent the output of the predictive function $p_k(g(f(\mathbf{x}_t)))$, at time $t+k$, at the spatial location (i, j) . We use a recursive context generation:

$$\begin{aligned} \mathbf{p}_{t+1} &= p_1(g(f(\mathbf{x}_t))) \\ &\dots \quad \dots \quad \dots \\ \mathbf{p}_{t+k} &= p_1(g(\mathbf{p}_{t+k-1})) \end{aligned}$$

¹ We drop the dependency on the parameters θ for the sake of clarity

Finally, a *regressor head* $c(\cdot; \theta_c)$ is placed on top of the GRU to perform a per-frame Action Unit intensity estimation. This head takes the hidden state at time \mathbf{h}_t , and after applying an average pooling operation, forwards the 256-dimensional frame representation to a simple block consisting of a Batch Normalization layer [19], and a linear layer that produces an N -dimensional output, where N is the number of the target AUs.

3.3 Unsupervised pre-training

It is worth noting that, even when working in a fully supervised manner, training the above model is a rather hard task. For the purposes of semi-supervised learning, we propose to use instead a self-supervised pre-training inspired by the contrastive learning [36], which allows us to make the network produce strong facial feature representations with no labels.

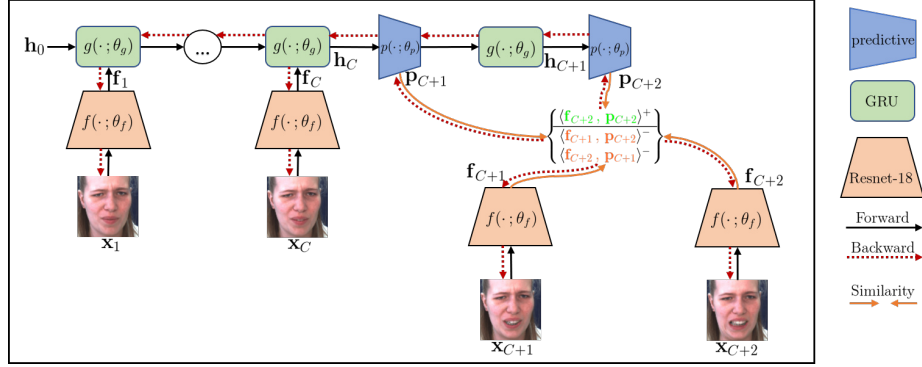


Fig. 3. Self-supervised learning through predictive coding. A predictive function uses the hidden state \mathbf{h}_C of the GRU, at time t_C , to predict the feature representation \mathbf{p}_{C+1} at the next temporal step t_{C+1} . This prediction is also used to obtain the next hidden state \mathbf{h}_{C+1} . The feature prediction is then repeated to obtain \mathbf{p}_{C+2} . The pair given by \mathbf{p}_{C+2} and \mathbf{f}_{C+2} (in green), is used as a positive sample, whereas all other mixed pairs are used as negatives (orange pairs). For the sake of clarity we illustrate only one positive pair. All other similar correspondences act also as positives during training, see Sec 3.3. We can then backpropagate (red lines) the similarity scores for the given pairs w.r.t. the contrastive loss defined in Eqn. 1, and learn a strong feature representation without labels.

i.e. we enforce the predictions to be conditioned not only on previous observations, but also on the previous predictions. Recall that $g(f(\mathbf{x}_t)) = g(f(\mathbf{x}_t), \mathbf{h}_{t-1})$, i.e. all frames before t are summarized in the context at time t .

The learning is accomplished through a Noise Contrastive Estimation [13, 12], where the goal is to classify real from noisy samples. Real samples are in practice formed by pairs $(\mathbf{f}_{t+k}^{(i,j)}, \mathbf{p}_{t+k}^{(i,j)})$, while noisy samples are formed by pairs $(\mathbf{f}_{t+k'}^{(i',j')}, \mathbf{p}_{t+k}^{(i,j)})$. In other words, noisy samples are formed by pairs composed of the feature representation at time t and spatial location (i, j) , and all the predictions taken from the same time position at different spatial locations, the predictions taken at different time steps, and even the predictions computed at different images in a given batch. For a given set of P predicted representations with size $H \times W$, the loss is formulated as:

$$\mathcal{L}_{nce} = - \sum_{k,i,j} \left[\log \frac{e^{\langle \mathbf{f}_{t+k}^{(i,j)}, \mathbf{p}_{t+k}^{(i,j)} \rangle}}{\sum_{k',i',j'} e^{\langle \mathbf{f}_{t+k'}^{(i',j')}, \mathbf{p}_{t+k}^{(i,j)} \rangle}} \right] \quad (1)$$

where $\langle \cdot, \cdot \rangle$ denotes the dot product, and is used as a similarity score between the feature representations. In Eqn. 1, $k \in \{1 \dots P\}$, and $(i, j) \in \{(1, 1) \dots (H, W)\}$. The loss within the brackets represents the typical cross-entropy objective used for classification, where the goal is to classify the positive pair among a set of $P \times H \times W$ classes (i.e. pairwise scores). When the set of negatives is enhanced with other images in the batch, the set of possible classes becomes $B \times P \times H \times W$,

with B the batch size. In our setting, $B = 20$, $P = 5$, and $W = H = 8$, as aforementioned, making the number of predictions be 6400. Note that in [36] the number of classes is dozens. This makes our predictive task much harder, making the representations be more locally distinct. The reason behind this approach relies on that we want the representation to be strong for a downstream task based on a *per-frame* classification. Minimizing Eqn. 1 leads to learning the weights θ_f and θ_g without requiring labeled videos. This procedure allows us to learn a strong representation with videos collected in-the-wild. A visual description is shown in Fig. 3.

3.4 Learning with partially labeled data

Now, we turn into how to train the network when only a few labeled facial images, randomly and discontinuously sampled, are available. In this paper, we propose a simple approach that consists of choosing random windows around a labeled frame, so that the position of the latter in the sequence varies each time it is queried for updating the network parameters. Using this approach, we apply back-propagation w.r.t. the output at the labeled frame. In other words, if only the frame \mathbf{x}_t is labeled within a sequence $\mathbf{X} = \{\mathbf{x}\}_{t=1}^T$, the classifier $c(\cdot; \theta_c)$ is updated only with a given labeled frame, and the feature extractor and temporal units, $f(\cdot; \theta_f)$ and $g(\cdot; \theta_g)$, are updated through back-propagation through time up to the labeled frame t . Denoting with \mathbf{y}_t the label for frame t , the loss is formulated as:

$$\mathcal{L}_{sup} = \|\mathbf{y}_t - c(g(f(\mathbf{x}_t)))\|^2, \quad (2)$$

where it is important to recall the dependency of $c(g(f(\mathbf{x}_t)))$ with all frames up to t . This way, the parameters θ_f and θ_g are updated with all the frames in the sequence up to t , and the parameters θ_c are updated only with the context given by g at time t and the corresponding labeled frame. This approach is illustrated in Fig. 4.

4 Experimental results

Data Throughout our experiments, we use three different datasets. For the self-supervised pretraining described in Section 3.3, we used the raw videos from the **Aff-Wild2** database [22, 25, 26, 24, 27, 23, 55] (i.e. with no annotations). We use 422 videos from Aff-Wild2 with around 1.2 million frames in total. In order to learn the feature representation using the predictive task, we used 351 videos for training, and 71 videos for validation. For the Action Unit intensity estimation, we use the **BP4D** and **DISFA** datasets. The **BP4D** database [58] is the main corpus of the **FERA2015** [45], and consists of videos of 41 subjects performing 8 different tasks, making it a total of 328 videos. For our experiments, we use the official train and validation subject-independent partitions, consisting of 21 and 20 subjects, respectively. The database contains around 140,000 frames, and is annotated with AU intensity for 5 AUs. In addition to the BP4D, we evaluate our method and perform a thorough ablation study on the **DISFA** database [32],

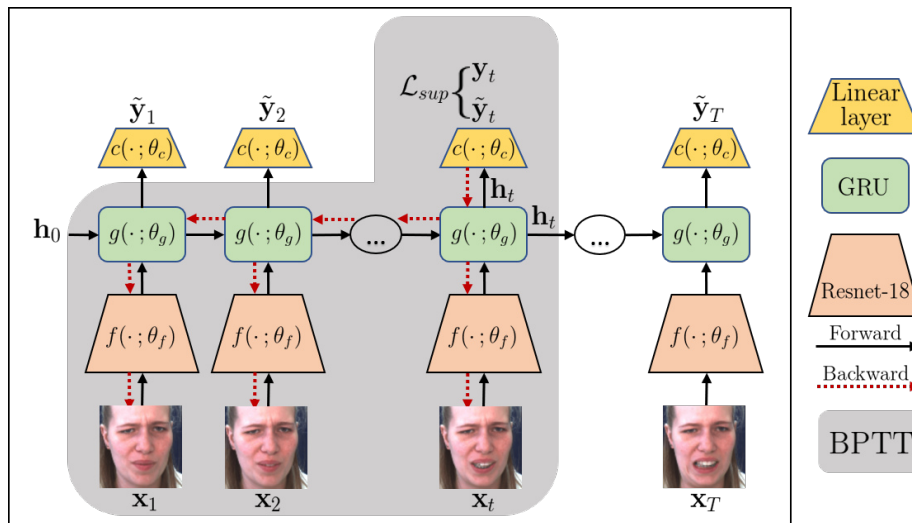


Fig. 4. Learning with partially labeled data: the network produces a structured output $\{\tilde{y}_t\}_{t=1}^T$, but only one frame y_t is labeled. We apply Backpropagation through time (BPTT) w.r.t. the labeled frame. While the labeled frames remain the same through the training process, their position in a queried sequence is randomly shifted (i.e. the number of frames before and after the labeled frame are randomly varied).

which includes 27 videos, each of a different subject, while performing computer tasks, for an average of ~ 4 minutes. It comprises around 130,000 frames annotated with the intensity of 12 AUs. To compare our method w.r.t. state-of-the-art results, we perform a subject independent three-fold cross-validation, where, for each fold, 18 subjects are used for training and 9 for testing. For our ablation studies, we use only one of the three folds (the same for all studies).

Frame selection Following existing works on semi-supervised Action Unit intensity estimation, we consider a labeled set of 2% of frames. However, contrary to existing works, we select the labeled frames *randomly*. To avoid the results of our ablation studies to be different due to the choice of data, we use the same subset of images to train both our models and those used for ablation studies. The number of training images for BP4D is 1498, whereas the number of images for each fold in DISFA is 1162. Fig. 5 shows the AU intensity distribution.

Evaluation metrics We use the Intra-Class Correlation (ICC(3,1) [42]), commonly used to rank methods in existing benchmarks. For an AU j with ground-truth labels $\{y_i^j\}_{i=1}^N$, and predictions $\{\hat{y}_i^j\}_{i=1}^N$, the ICC score is defined as $ICC^j = \frac{W^j - S^j}{W^j + S^j}$, with $W^j = \frac{1}{N} \sum_i \left((y_i^j - \hat{y}^j)^2 + (\tilde{y}_i^j - \hat{y}^j)^2 \right)$, $S^j = \sum_i (y_i^j - \tilde{y}_i^j)^2$, and $\hat{y}^j = \frac{1}{2N} \sum_i (y_i^j + \tilde{y}_i^j)$. In addition, we report the Mean Absolute Error (MAE).




Fig. 5. AU intensity distribution for DISFA (left) and BP4D (right).

Set up We use the PyTorch automatic differentiation package in all our experiments [37]. We use the publicly available S³FD face detector of [57] to detect a facial bounding box in each video. The images are then tightly cropped around the center and resized to 128×128 . Some augmentation is applied to the images for both the self-supervised training and the final stage (see below).

Self-supervised training The training is composed of 33509 sequences of 15 frames each, sampled with a stride of 2, and extracted from Aff-Wild2. To avoid discontinuities, sequences where a face is not detected are discarded. The number of valid test sequences to validate the performance of the predictor is 7164. To minimize non-facial content we tightly crop around the center of the bounding box. We apply a set of heavy augmentations during training, including uniform rotation (± 20 deg.), uniform scaling ($\pm 10\%$), random flipping, random jitter, and random hue, contrast, saturation and brightness. The number of context frames is set to 10, and the set of predicted frames is 5. We use the Adam optimizer [21], with an initial learning rate of 10^{-3} and weight decay of 10^{-5} . The model is trained for 300 epochs using 8 GPUs, each having a batch size of 20 sequences. The training takes approximately 1 day to be completed. We evaluate the capacity of the network to pick up the right pair by measuring the Top- n accuracy on the validation set, with $n = 1, 3, 5$. In other words, Top- n shows the percentage of samples in the validation set where the score corresponding to the positive pair was among the highest n scores. The results are shown in the second row of Table 1. As an ablation study, we also used a subset of DISFA to perform the self-supervised training. See Section 4.1 for further details. The Top- n accuracy for the DISFA is shown in the first row of Table 1.

Semi-supervised training To train the Action Unit intensity regressor, we use learning rate 10^{-5} with weight decay 10^{-5} . We use Adam with $\beta = (0.5, 0.999)$. For the model trained from scratch, we use a Kaiming initialization [16]. We observed that the model is sensitive to the initialization. When using

Accuracy	Top-1	Top-3	Top-5
DISFA	.165	.467	.695
Aff-Wild2	.303	.651	.781

Table 1. Accuracy on the validation set of DISFA and Aff-Wild2 for the Contrastive Predictive task. See Sec 4.1 for details

normally distributed random weights, the training was unstable, and the performance was very poor.

4.1 Ablation study

We first perform an ablation study to validate the effectiveness of our proposed approach w.r.t. different alternatives. To this end, we use the most challenging fold from DISFA to train and validate each of the models, i.e. we used 18 subjects to train and 9 to test. The best performing model for each method is chosen. In particular, we illustrate the results of the following options in Table 2 and Table 3.

- **R18 - 2%** We train a simple ResNet-18 on the same 2% of the data, by minimizing the \mathcal{L}_2 loss between the available labels and the predictions on the corresponding images. We use the same augmentation as mentioned above.
- **R18 - Sup.** We train the same ResNet-18 on the full dataset (58140 images).
- **R18+GRU-scratch - 2%** We train the whole pipeline without the self-supervised training described in Section 3.3, i.e. we train the whole model from scratch.
- **R18+GRU-scratch Sup.** We also trained the whole pipeline without the self-supervised training in a fully supervised manner.
- **R18+GRU-scratch Pseudo-GT** In this setting, we used a different approach to semi-supervised learning, that of pseudo-labeling. In particular, we used the R18 - 2% model to generate the labels for the training set, and used these labels to train our model. In addition, the weights of the R18-2% network are also used to initialize the R18+GRU pipeline.
- **Ours(*DISFA) - 2%** In this setting, we used the same training partition to perform the self-supervised training described in Section 3.3, i.e. we used the 18 training videos from DISFA. Then, we initialized our network with the generated weights, and trained using only the aforementioned 2% of labels.
- **Ours - 2%** This setting corresponds to using our proposed approach to learn the model with 2% of the data.
- **Ours - Sup.** We finally evaluate the performance of the model when the whole training set is available. This serves as an upper bound in the performance of the proposed approach.

Discussion The results of all these models are shown in Table 2 and Table 3. From these results, we can make the following observations:

- *1. Accuracy of self-supervised learning* The results shown in Table 1 indicate that the performance in the predictive task is superior when the network is trained on Aff-Wild2 than when it is trained on DISFA. We attribute this difference to the fact that DISFA has less variability than Aff-Wild2, mainly due to the recording conditions. In addition, the training set for DISFA is composed of barely 18 videos, and hence the number of negative samples are highly likely to include segments that are too similar. In this scenario, the negative samples become almost indistinguishable from the positive ones, making the predictive task harder to learn.

Dataset AU	DISFA - Only 2% of frames are annotated												Avg.
	1	2	4	5	6	9	12	15	17	20	25	26	
R18 2%	[.144]	-.014	.375	[.129]	.499	[.366]	.737	.208	.355	[.210]	.901	.458	.364
⊂ R18+GRU-scratch 2%	-.032	-.126	.340	.009	.393	.043	.650	.077	.048	.004	.841	.434	.223
⊂ Ours(*Disfa)-2%	.084	[.070]	.489	.033	.481	.225	[.772]	[.337]	.304	.127	.853	[.578]	.363
⊂ Ours(*Affwild2)-2%	-.010	-.022	[.657]	.068	[.566]	.358	.737	.291	[.366]	.109	[.944]	.537	[.383]
R18 2%	[.209]	[.228]	1.022	[.042]	.339	.311	.333	[.137]	.306	.191	.307	.474	.325
⊂ R18+GRU-scratch 2%	.555	.512	.984	.076	.442	.397	.542	.178	.252	[.150]	.546	.436	.423
⊂ Ours(*Disfa)-2%	.238	.236	[.683]	.098	.329	.287	[.326]	.174	[.305]	.217	.428	.389	[.309]
⊂ Ours(*Affwild2)-2%	.392	.463	.754	.129	[.304]	[.277]	.360	.189	.386	.201	[.260]	[.298]	.334

Table 2. Ablation study on DISFA in a scenario where only 2% of frames are annotated. R18-2% refers to a simple ResNet-18 trained with 2% of the data. R18+GRU-scratch-2% refers to the method described in Section 3.4, without the self-supervised learning. Ours(*Disfa) refers to our method, with the self-supervised learning stage being done on the DISFA dataset. Ours(*Affwild2)-2% refers to our method, trained on 2% of the data. **[bold]** indicates best performance

- *2. Importance of pre-text training set* In addition to the lower performance on the predictive task, it is worth noting that the learned representation when using few videos is much weaker than that learned from videos collected in the wild. This is illustrated in the performance that the network yields in the task of facial Action Unit intensity estimation. It is important to remark that collecting videos in-the-wild that *do not* require labeling is nowadays cheaper than annotating all the frames with 12 Action Units even for a small number of videos.
- *3. Influence of self-supervised learning* We observe that, whether on DISFA or on Aff-Wild2, the results on the downstream task are remarkably better than training the network from scratch. We observe this not only for our case of interest where few labels are available, but also when training the network in a fully supervised manner. We also observed that, when trained from scratch, the network is certainly sensitive to the initialization. This has been observed in [16], that indicates that a poor initialization with deep networks can lead to vanishing gradients, making the training process unstable. In our experiments, we observe that training the Resnet-18 in a supervised manner without any temporal modeling yields better results than training the whole pipeline, also in a supervised manner. We can see that our method is effective both when training with few labels and when training in a supervised way.
- *4. Influence of pseudo-labeling* We observe that this technique can be powerful enough, yielding an ICC score similar to that given by training the network from scratch in a fully supervised manner.

4.2 Comparison with state-of-the-art

We now show the results of our method w.r.t. reported state of the art methods on weakly supervised learning for facial Action Unit intensity estimation. We report our results for both the BP4D case and for the three-fold cross validation performed in DISFA, so as to make our results comparable to existing

Dataset AU	DISFA - All frames are annotated												Avg.
	1	2	4	5	6	9	12	15	17	20	25	26	
R18 Sup.	.215	.030	.470	[-339]	[-583]	.382	.790	[-387]	.476	.383	.887	.565	.459
R18+GRU-scratch Sup.	.047	-.056	.366	.097	.455	.145	.775	.089	.313	.126	.812	.501	.306
R18+GRU-Pseudo-GT	.096	.017	.469	.115	.470	.227	.653	.186	.317	.257	.578	.250	.303
Ours (*Affwild2)-Sup.	[-291]	[-107]	[-495]	.281	.499	[-424]	[-798]	.327	[-513]	[-473]	[-908]	[-648]	[-480]
R18 Sup.	.216	.222	1.002	.048	[-302]	.302	.314	[-117]	[-200]	.134	.321	[-346]	.294
R18+GRU-scratch Sup.	.365	.341	1.022	.113	.364	.553	.352	.204	.444	.259	.601	.549	.431
R18+GRU-Pseudo-GT	.225	.374	.816	.062	.640	.282	.530	.296	.844	.185	1.127	.631	.501
Ours (*Affwild2)-Sup.	[-161]	[-200]	[-815]	[-043]	.334	[-273]	[-292]	.127	.215	[-108]	[-291]	.362	[-268]

Table 3. Ablation study on DISFA where a fully supervised setting is applied. R18 Sup. corresponds to a Resnet-18 trained in a supervised manner. R18+GRU-scratch-Sup. refers to the supervised training. R18+GRU-scratch-PseudoGT refers to a fully supervised training using the pseudo-labels produced by the R18-2% shown in Table 2. Ours(*Affwild2)-Sup. refers to our method, trained with all the labels. **[bold]** indicates best performance.

works. We show the results of our method for BP4D in Table 4, and the results for DISFA in Table 5. We compare our method with the most recent works reporting Action Unit intensity estimation with partially labeled data. In particular, we compare the performance of our method with that of **KBSS** [60], **KJRE** [62], and **CFLF** [61]. Importantly, the three methods require the labeling of key frames, mainly due to the fact that one of the components for weakly supervised learning relies on assuming that intermediate frames between keys follow some monotonic behavior. Also, both the KBSS and CFLF work under the basis of a 1% of labeled frames. However, these methods, reportedly, use a different percentage of frames *per Action Unit*. Different from these methods, our proposed approach works with a randomly chosen subset of data. As shown in both Table 4 and Table 5, our method yields competitive results, and surpasses the three aforementioned methods in terms of average ICC.

	Dataset AU	BP4D					Avg.
		6	10	12	14	17	
ICC	KBSS [60]*	.760	.725	.840	.445	.454	.645
	KJRE [62]* 6%	.710	.610	[-.870]	.390	.420	.600
	CFLF [61]*	[-.766]	.703	.827	.411	[-.600]	.661
	Ours 2%	[-.766]	[-.749]	.857	[-.475]	.553	[-.680]
MAE	KBSS [60]*	.738	[-.773]	.694	.990	.895	.818
	KJRE [62]* 6%	.820	.950	[-.640]	1.080	.850	.870
	CFLF [61]*	[-.624]	.830	.694	1.000	[-.626]	[-.741]
	Ours 2%	.645	.913	.826	[-.979]	.628	.798

Table 4. Intensity estimation results on BP4D. (*) Indicates results taken from reference. **[bold]** indicates best performance.

Dataset AU	DISFA												Avg.
	1	2	4	5	6	9	12	15	17	20	25	26	
KBSS [60]*	.136	.116	.480	.169	.433	.353	.710	.154	.248	.085	.778	.536	.350
UC KJRE [62]* 6%	.270	[.350]	.250	.330	.510	.310	.670	.140	.170	.200	.740	.250	.350
CLFL [61]*	.263	.194	.459	[.354]	.516	[.356]	.707	[.183]	[.340]	[.206]	.811	.510	.408
Ours 2%	[.327]	.328	[.645]	-.024	[.601]	.335	[.783]	.181	.243	.078	[.882]	[.578]	[.413]
KBSS [60]*	.532	.489	.818	.237	.389	.375	.434	.321	.497	.355	.613	.440	.458
UC KJRE [62]* 6%	1.020	.920	1.860	.700	.790	.870	.770	.600	.800	.720	.960	.940	.910
CLFL [61]*	[.326]	[.280]	[.605]	[.126]	[.350]	[.275]	[.425]	[.180]	[.290]	[.164]	.530	.398	[.329]
Ours 2%	.430	.358	.653	.194	.381	.370	.457	.247	.376	.212	[.446]	[.387]	.376

Table 5. Intensity estimation results on DISFA. (*) Indicates results taken from reference. **[bold]** indicates best performance.

5 Conclusion

In this paper, we proposed a novel approach to semi-supervised training of facial Action Unit intensity estimation, that is capable of delivering competing results with as little as 2% of annotated frames. To this end, we proposed a self-supervised learning approach that can capture strong semantic representations of the face, that can later be used to train models in a semi-supervised way. Our approach surpasses existing works on semi-supervised learning of Action Unit intensity estimation. We also demonstrated that our approach, when used in a fully supervised manner, largely outperforms a model trained from scratch, thus demonstrating that our approach is also valid for supervised learning. The experimental evaluation proved the effectiveness of our method, through several ablation studies.

References

- Bachman, P., Hjelm, R.D., Buchwalter, W.: Learning representations by maximizing mutual information across views. In: NeurIPS (2019)
- Caron, M., Bojanowski, P., Joulin, A., Douze, M.: Deep clustering for unsupervised learning of visual features. In: ECCV (2018)
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y.: Learning phrase representations using rnn encoder-decoder for statistical machine translation. EMNLP (2014)
- Chu, W., De la Torre, F., Cohn, J.F.: Learning spatial and temporal cues for multi-label facial action unit detection. In: FG (2017)
- Chu, W.S., la Torre, F.D., Cohn, J.F.: Learning facial action units with spatiotemporal cues and multi-label sampling. Image and Vision Computing (2019)
- Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: ICCV (2015)
- Ekman, P., Friesen, W., Hager, J.: Facial action coding system. In: A Human Face (2002)
- Eleftheriadis, S., Rudovic, O., Deisenroth, M.P., Pantic, M.: Variational gaussian process auto-encoder for ordinal prediction of facial action units. In: ACCV (2016)
- Ertugrul, I.O., Cohn, J.F., Jeni, L.A., Zhang, Z., Yin, L., Ji, Q.: Crossing domains for au coding: Perspectives, approaches, and measures. IEEE transactions on biometrics, behavior, and identity science **2**(2), 158–171 (2020)

10. Ertugrul, I.O., Jeni, L.A., Cohn, J.F.: Pattnet: Patch-attentive deep network for action unit detection. In: BMVC (2019)
11. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations (2018)
12. Gutmann, M., Hyvärinen, A.: Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In: AISTATS (2010)
13. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: CVPR (2006)
14. Han, T., Xie, W., Zisserman, A.: Video representation learning by dense predictive coding. In: CVPR - Workshops (2019)
15. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. arXiv preprint arXiv:1911.05722 (2019)
16. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: ICCV (2015)
17. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
18. Hénaff, O.J., Razavi, A., Doersch, C., Eslami, S., Oord, A.v.d.: Data-efficient image recognition with contrastive predictive coding. arXiv preprint arXiv:1905.09272 (2019)
19. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: ICML (2015)
20. Jaiswal, S., Valstar, M.: Deep learning the dynamic appearance and shape of facial action units. In: Winter Conference on Applications of Computer Vision (2016)
21. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. International Conference on Learning Representations (2014)
22. Kollias, D., Nicolaou, M.A., Kotsia, I., Zhao, G., Zafeiriou, S.: Recognition of affect in the wild using deep neural networks. In: CVPR - Workshops. pp. 1972–1979. IEEE (2017)
23. Kollias, D., Schulc, A., Hajiyev, E., Zafeiriou, S.: Analysing affective behavior in the first abaw 2020 competition. arXiv preprint arXiv:2001.11409 (2020)
24. Kollias, D., Tzirakis, P., Nicolaou, M.A., Papaioannou, A., Zhao, G., Schuller, B., Kotsia, I., Zafeiriou, S.: Deep affect prediction in-the-wild: Aff-wild database and challenge, deep architectures, and beyond. IJCV pp. 1–23 (2019)
25. Kollias, D., Zafeiriou, S.: Aff-wild2: Extending the aff-wild database for affect recognition. arXiv preprint arXiv:1811.07770 (2018)
26. Kollias, D., Zafeiriou, S.: A multi-task learning & generation framework: Valence-arousal, action units & primary expressions. arXiv preprint arXiv:1811.07771 (2018)
27. Kollias, D., Zafeiriou, S.: Expression, affect, action unit recognition: Aff-wild2, multi-task learning and arface. arXiv preprint arXiv:1910.04855 (2019)
28. Li, W., Abtahi, F., Zhu, Z.: Action unit detection with region adaptation, multi-labeling learning and optimal temporal fusing. In: CVPR (2017)
29. Li, W., Abtahi, F., Zhu, Z., Yin, L.: Eac-net: A region-based deep enhancing and cropping approach for facial action unit detection. In: FG (2017)
30. Li, Y., Zeng, J., Shan, S., Chen, X.: Self-supervised representation learning from videos for facial action unit detection. In: CVPR (2019)
31. Martinez, B., Valstar, M.F., Jiang, B., Pantic, M.: Automatic analysis of facial actions: A survey. IEEE transactions on affective computing (2017)
32. Mavadati, S.M., Mahoor, M.H., Bartlett, K., Trinh, P., Cohn, J.F.: Disfa: A spontaneous facial action intensity database. IEEE-TAC (2013)

33. Ming, Z., Bugeau, A., Rouas, J., Shochi, T.: Facial action units intensity estimation by the fusion of features with multi-kernel support vector machine. In: FG (2015)
34. Misra, I., Zitnick, C.L., Hebert, M.: Shuffle and learn: unsupervised learning using temporal order verification. In: ECCV (2016)
35. Ntinou, I., Sanchez, E., Bulat, A., Valstar, M., Tzimiropoulos, G.: A transfer learning approach to heatmap regression for action unit intensity estimation. arXiv preprint arXiv:2004.06657 (2020)
36. van den Oord, A., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
37. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch. In: Autodiff workshop - NeurIPS (2017)
38. Rudovic, O., Pavlovic, V., Pantic, M.: Automatic pain intensity estimation with heteroscedastic conditional ordinal random fields. In: Advances in Visual Computing (2013)
39. Rudovic, O., Pavlovic, V., Pantic, M.: Context-sensitive conditional ordinal random fields for facial action intensity estimation. In: ICCV - Workshops (2013)
40. Rudovic, O., Pavlovic, V., Pantic, M.: Context-sensitive dynamic ordinal regression for intensity estimation of facial action units. IEEE-TPAMI (2015)
41. Sanchez, E., Tzimiropoulos, G., Valstar, M.: Joint action unit localisation and intensity estimation through heatmap regression. In: BMVC (2018)
42. Shrout, P., Fleiss, J.: Intraclass correlations: uses in assessing rater reliability. Psychological bulletin (1979)
43. Tian, Y., Krishnan, D., Isola, P.: Contrastive multiview coding. arXiv preprint arXiv:1906.05849 (2019)
44. Tran, D.L., Walecki, R., Rudovic, O., Eleftheriadis, S., Schuller, B., Pantic, M.: Deepcoder: Semi-parametric variational autoencoders for automatic facial action coding. In: ICCV (2017)
45. Valstar, M.F., Almaev, T., Girard, J.M., McKeown, G., Mehu, M., Yin, L., Pantic, M., Cohn, J.F.: Fera 2015 - second facial expression recognition and analysis challenge. In: FG (2015)
46. Vincent, P., Larochelle, H., Bengio, Y., Manzagol, P.A.: Extracting and composing robust features with denoising autoencoders. In: ICML. pp. 1096–1103 (2008)
47. Walecki, R., Rudovic, O., Pavlovic, V., Schuller, B., M.Pantic: Deep structured learning for facial action unit intensity estimation. In: CVPR (2017)
48. Wang, S., Peng, G.: Weakly supervised dual learning for facial action unit recognition. IEEE Transactions on Multimedia **21**(12) (2019)
49. Wang, S., Pan, B., Wu, S., Ji, Q.: Deep facial action unit recognition and intensity estimation from partially labelled data. IEEE Transactions on Affective Computing (2019)
50. Werbos, P.J.: Backpropagation through time: what it does and how to do it. Proceedings of the IEEE **78**(10), 1550–1560 (1990)
51. Wu, Y., Ji, Q.: Constrained joint cascade regression framework for simultaneous facial action unit recognition and facial landmark detection. In: CVPR (2016)
52. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: CVPR (2018)
53. Yang, L., Ertugrul, I.O., Cohn, J.F., Hammal, Z., Jiang, D., Sahli, H.: Facs3d-net: 3d convolution based spatiotemporal representation for action unit detection. In: ACII (2019)
54. Ye, M., Zhang, X., Yuen, P.C., Chang, S.F.: Unsupervised embedding learning via invariant and spreading instance feature. In: CVPR (2019)

55. Zafeiriou, S., Kollias, D., Nicolaou, M.A., Papaioannou, A., Zhao, G., Kotsia, I.: Aff-wild: Valence and arousal ‘in-the-wild’ challenge. In: CVPR - Workshops. IEEE (2017)
56. Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: ECCV (2016)
57. Zhang, S., Zhu, X., Lei, Z., Shi, H., Wang, X., Li, S.Z.: S3fd: Single shot scale-invariant face detector. In: ICCV (2017)
58. Zhang, X., Yin, L., Cohn, J.F., Canavan, S., Reale, M., Horowitz, A., Liu, P., Girard, J.M.: Bp4d-spontaneous: a high-resolution spontaneous 3d dynamic facial expression database. *Image and Vision Computing* (2014)
59. Zhang, Y., Dong, W., Hu, B.G., Ji, Q.: Classifier learning with prior probabilities for facial action unit recognition. In: CVPR (2018)
60. Zhang, Y., Dong, W., Hu, B.G., Ji, Q.: Weakly-supervised deep convolutional neural network learning for facial action unit intensity estimation. In: CVPR (2018)
61. Zhang, Y., Jiang, H., Wu, B., Fan, Y., Ji, Q.: Context-aware feature and label fusion for facial action unit intensity estimation with partially labeled data. In: ICCV (2019)
62. Zhang, Y., Wu, B., Dong, W., Li, Z., Liu, W., Hu, B.G., Ji, Q.: Joint representation and estimator learning for facial action unit intensity estimation. In: CVPR (2019)
63. Zhang, Y., Zhao, R., Dong, W., Hu, B.G., Ji, Q.: Bilateral ordinal relevance multi-instance regression for facial action unit intensity estimation. In: CVPR (2018)
64. Zhao, K., Chu, W.S., Zhang, H.: Deep region and multi-label learning for facial action unit detection. In: CVPR (2016)
65. Zhao, R., Gan, Q., Wang, S., Ji, Q.: Facial expression intensity estimation using ordinal information. In: CVPR (2016)