

Expression, Affect, Action Unit Recognition: Aff-Wild2, Multi-Task Learning and ArcFace

Dimitrios Kollias^{1,2}
dimitrios.kollias15@imperial.ac.uk

¹ Department of Computing
Imperial College London
London, UK

Stefanos Zafeiriou^{1,2}
s.zafeiriou@imperial.ac.uk

² FaceSoft.io

Abstract

Affective computing has been largely limited in terms of available data resources. The need to collect and annotate diverse in-the-wild datasets has become apparent with the rise of deep learning models, as the default approach to address any computer vision task. Some in-the-wild databases have been recently proposed. However: i) their size is small, ii) they are not audiovisual, iii) only a small part is manually annotated, iv) they contain a small number of subjects, or v) they are not annotated for all main behavior tasks (valence-arousal estimation, action unit detection and basic expression classification). To address these, we substantially extend the largest available in-the-wild database (Aff-Wild) to study continuous emotions such as valence and arousal. Furthermore, we annotate parts of the database with basic expressions and action units. As a consequence, for the first time, this allows the joint study of all three types of behavior states. We call this database Aff-Wild2. We conduct extensive experiments with CNN and CNN-RNN architectures that use visual and audio modalities; these networks are trained on Aff-Wild2 and their performance is then evaluated on 10 publicly available emotion databases. We show that the networks achieve state-of-the-art performance for the emotion recognition tasks. Additionally, we adapt the ArcFace loss function in the emotion recognition context and use it for training two new networks on Aff-Wild2 and then re-train them in a variety of diverse expression recognition databases. The networks are shown to improve the existing state-of-the-art. The database, emotion recognition models and source code are available at <http://ibug.doc.ic.ac.uk/resources/aff-wild2>.

1 Introduction

Until recently affective computing has been mostly studied in controlled settings of the environments [14, 25, 39, 47, 48], with limited amount of participants [3, 26, 31, 33, 40], using pre-defined scenarios that users have to follow, depicting posed expressions [10, 28]. However, with the development of large and diverse datasets in the field of computer vision (and the accompanying performance gains), it has

become apparent that the diversity of human participants and spontaneous expressions have to become the prerogatives in deployment of the affective computing models in practice. Large datasets with in-the-wild settings have been recently collected to study facial expression (Expr) analysis [6, 7], facial action units (AUs) [10] and continuous emotions of valence and arousal (VA) [63, 65] in-the-wild.

In [29] a static in-the-wild database (AffectNet) has been created that contains VA annotations for 1M images. However, from those images, only around 450K are manually annotated and from those only 350K are valid faces. This number is moderate for training deep neural networks (DNNs). Also this database is static, meaning that it contains only images and no video/audio. It is worth to mention that this database also contains annotations for the seven basic expressions plus the contempt class for 290K images. In [9] a static in-the-wild database (Emotionet) has been created that contains AU annotations for 1M images. However, from those images, only 50K are manually annotated. Half of those consist the validation set and the other half the test set. Again, these sets have a small size to adequately train DNNs and generalize on other databases. Emotionet database also contains annotations for 6 basic and 10 compound emotion categories. Nevertheless, their total size is less than 3K and the classes are heavily imbalanced, making it impossible to train DNNs. In [20, 50], the authors have developed the largest existing audiovisual (A/V) in-the-wild database annotated in terms of VA for around 1.25M. All these annotations are manual. However this database contains annotations only for VA and the number of subjects in the videos is moderate (298 subjects in total).

Up to the present, there is no database that contains annotations for all main behavior tasks (VA estimation, AU detection, Expr classification). Also, most of the existing databases do not contain sufficiently large numbers of annotated samples for effectively training DNNs. The first contribution in this paper is the creation of a new dataset that contains 260 videos with around 1.4M frames, annotated for VA. We merge this dataset with Aff-Wild (since this database also contains annotated videos), generating the so-called Aff-Wild2 database. Next, we annotate parts of Aff-Wild2 with AUs and seven basic expression labels, creating about 398K and 403K AU and Expr annotations, respectively. To the best of our knowledge, Aff-Wild2 is the first large scale in-the-wild database containing annotations for all 3 main behavior tasks. It is also the first audiovisual database with annotations for AUs. All AU annotated databases do not contain audio, but only images or videos.

Next, we conduct multi-task experiments on this database, for emotion recognition. Many questions arise hereafter: how can we combine these three tasks? what loss function should we use? The first apparent answer is to use a loss function equal to the sum of the loss functions of each task. The binary cross entropy loss is used for AU detection. The MSE and Concordance Correlation Coefficient (CCC) losses are used for VA estimation. The standard loss for expression classification is the categorical cross entropy. The second contribution of the paper is the development of multi-task CNNs, multi-task CNN-RNNs and multi-modal, multi-task CNN-RNNs, which are trained on Aff-Wild2 and then applied to 10 publicly available databases (including the Aff-Wild one). The results are very promising, beating the state-of-the-art on emotion recognition in these databases; exceptions are two databases annotated for Expr Recognition. In one of them, a best performing network used a locality preserving loss function [22]. This, as well as the recent tendency to develop elaborate loss functions for specific tasks [18], has led us to search

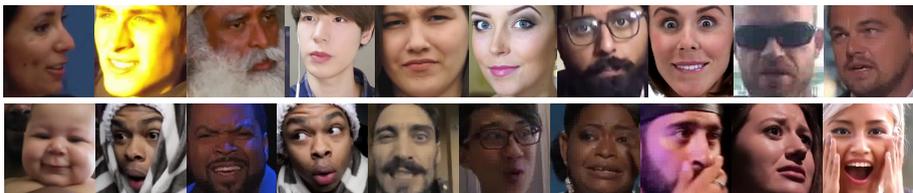


Figure 1: Frames of Aff-Wild2, showing subjects of different ethnicities, age groups, emotional states, head poses, illumination conditions and occlusions

for a better loss function than the categorical cross entropy.

In fact, in the related face recognition field, it has been shown [24, 24] that categorical cross entropy loss is insufficient to acquire discriminating power for face classification. Several loss functions have been proposed for maximizing inter-class and minimizing intra-class variance. [9, 15] propose multi-loss learning to increase feature discriminating power. These, require thorough mining of pair/triplet samples, which is a time-consuming procedure. [24] projects the original Euclidean space of features to an angular space, introducing an angular margin for larger inter-class variance. [43] directly adds a cosine margin penalty to the target logit, showing better performance than [24]. [8] further improved the discriminative power of face recognition models, stabilising the training process. As these losses boosted face recognition models performance, in this work, we choose to adopt the ArcFace loss [8] and adapt it for emotion recognition. To the best of our knowledge, this is the first time that such a loss designed for face recognition, is used in the context of emotion recognition. Our final contribution in this paper is the design of 2 networks trained with the ArcFace loss. After training them on Aff-Wild2, we re-trained them on each of the examined databases. Our results outperformed all state-of-the-art networks, illustrating: i) the richness of Aff-Wild2 (providing it with the ability to be used as robust prior for network pre-training) and ii) that the ArcFace loss can be used in the emotion recognition field, yielding state-of-the-art results. In fact, this is the very first proof of the effectiveness of additive angular margin in emotion recognition.

2 The Aff-Wild2 database

Aff-Wild2 is described next, presenting the new collected dataset and its properties, the generated partition sets, their distributions and the annotation procedure.

Collected dataset and properties We extend the Aff-Wild database [20, 50], by collecting a new dataset consisting of 260 YouTube videos, with 1,413,000 frames and a total length of 13 hours and 5 minutes. The videos have been collected using the Youtube video sharing website. All of the collected videos are in MP4 format, with a frame rate of 30, provided under the CC licence. Keywords for retrieving the videos were selected from the 2-D Emotion Wheel, shown in Figure 2. The new videos have wide range in subjects’: age (from babies to elderly people); ethnicity (caucasian/hispanic/latino/asian/black/african american); profession (e.g. actors, athletes, politicians, journalists); head pose; illumination conditions; occlusions; emotions. Figure 1 shows frames of Aff-Wild2 verifying the above described ranges.

These videos show subjects who: react on a surprise, on something that brings them happiness or fulfillment, on flirting or rejection, on important political issues, on funny or mean tweets; are stand-up comedians; give a really interesting speech in ceremonies; are taking an oral exam; are giving lectures on depression, or other serious disorders; are performing passive, boring, apathetic, intense activities, etc.

Four experts annotated the new dataset in terms of valence and arousal, as in the case of Aff-Wild. We then concatenated the Aff-Wild database with the new dataset, forming Aff-Wild2. In total, Aff-Wild2 consists of 558 videos with 2,786,201 frames, showing both subtle and extreme human behaviours in real-world settings. The total number of subjects is 458; 279 of which are males and 179 females.

Two more tasks were implemented, in which we annotated parts of Aff-Wild2 with AUs and Exprs. In the first, three very experienced annotators annotated 63 videos, with 397,800 frames and a total length of 3 hours and 41 mins, in terms of AUs 1,2,4,6,12,15,20,25 - described in Figure 2. These videos contain 31 male and 31 female subjects. In the second, three experts annotated 84 videos consisting of 403,758 frames, with a total length of 3 hours and 45 mins, in terms of the 7 basic expressions. The videos show 42 male and 42 female subjects. Consequently, Aff-Wild2 contains 3 datasets (VA, AU, Expr); each contains annotations for a respective behavior task (preliminary work regarding the 3 sets can be found in [16, 17]). Table 1 summarizes the attributes and properties of the three annotated sets of Aff-Wild2.

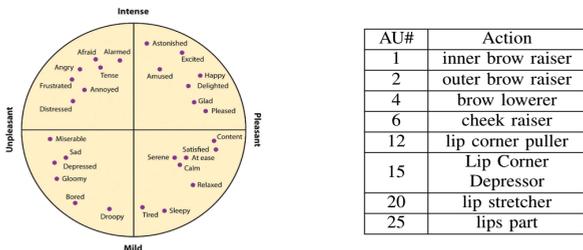


Figure 2: The 2D Emotion Wheel (left); the AUs annotated in Aff-Wild2 (right)

Table 1: General Attributes of Aff-Wild2; in the VA set, top row refers to the new dataset, while bottom row refers to Aff-Wild

Aff-Wild2	# frames	# videos	# annotators	Video Length	Mean Resolution
VA set	1,413,000	260	4	0.03 – 26.22 mins	1450 × 900
	1,373,201	298	8	0.10 – 14.47 mins	607 × 359
AU set	397,800	63	3	0.03 – 26.22 mins	1500 × 900
Expr set	403,758	84	3	0.04 – 26.22 mins	1350 × 800

Partition Sets and Distributions Each set (VA, AU, Expr) is split into three subsets: training, validation and test. Partitioning is done in a subject independent manner, in the sense that a person can appear only in one of those three subsets. In the VA set, the resulting training, validation and test subsets consist of 350, 70 and 138 videos respectively. In the AU set, the respective subsets consist of 42, 7 and 14 videos respectively. In the Expr set, the corresponding subsets consist of 51, 11 and 22 videos respectively. Figure 3 shows the 2D VA histogram of the new dataset, which was added to Aff-Wild. Figure 3 shows the distribution of the seven emotion categories in Aff-Wild2. Table 2 shows the distribution of the activated AUs. We

note that the Expr Set of images can be extended to also contain AU annotations, according to Table 1 of [10]. However, this is out of the scope of the current paper.

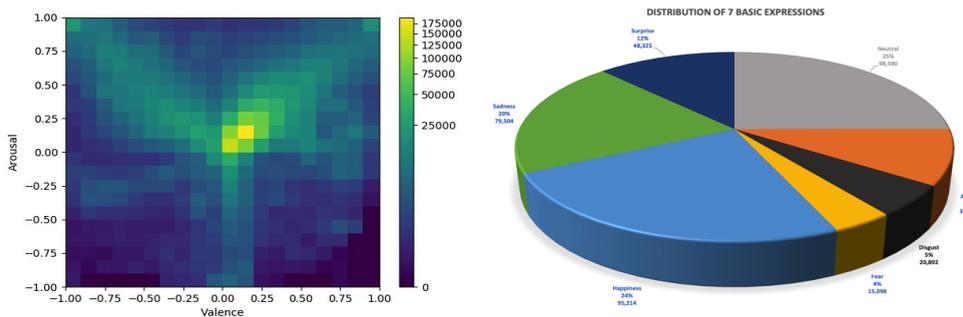


Figure 3: 2D VA Histogram of the new data added to Aff-Wild (left); Histogram of the seven basic expressions in Aff-Wild2 (right)

Table 2: Distribution of AU annotations in Aff-Wild2

Action Unit #	AU 1	AU 2	AU 4	AU 6	AU 12	AU 15	AU 20	AU 25
Total Number of Activated AUs	86,677 43.9 %	4,166 2.1 %	56,327 28.5 %	25,226 12.8 %	35,675 18.1 %	3,340 1.7 %	5,695 2.9 %	9,048 4.6 %

Annotation Four experts have performed the VA set annotation, using the method proposed in [6]. Valence and arousal values range continuously in $[-1,1]$. The final label values are the mean of those four annotations. The mean inter-annotation correlation is 0.63 for valence and 0.60 for arousal. For the AU set, three experts have performed the annotation. For the Expr set, three more experts performed the annotation. In both cases, agreement between the annotators has not always been 100%. We only kept the annotations, on which all experts agree.

3 Proposed Methods

Two pre-processing steps, on the visual and audio modalities, have been applied to generate the input data for DNN based emotion analysis. We developed the following deep network architectures for emotion recognition: i) CNN, single-/multi-task; ii) CNN-RNN multi-task; iii) CNN-RNN multi-modal (A/V) and multi-task; iv) new ArcFace networks with respective loss function, as described below.

Visual Modality Pre-Processing The SSH detector [30] based on the ResNet and trained on the WiderFace dataset [46] was used to extract face bounding boxes from all images. Also, 5 facial landmarks (two eyes, nose and two mouth corners) were extracted and used to perform similarity transformation (for face alignment). After that we obtain the cropped faces which are then resized to dimension $96 \times 96 \times 3$. The pixel intensities are normalized to take values in $[-1,1]$.

Audio Modality Pre-Processing The audio signal (mono) is sampled at 44,100Hz. Then spectrograms are extracted; spectrogram frames are computed over a 33ms

window with 11ms overlap. The resulting intensity values are normalized in $[-1,1]$ to be consistent with the visual modality.

CNN Single- & Multi-Task We employ 3 state-of-the-art networks, SphereFace-20 [24], VGGFace [52], and Inception ResNet [57] (denoted as Inc.ResNet). We train these networks to perform one behavior task (VA estimation, AU detection, or Expr classification), or jointly perform all 3 tasks. We call the multi-task VGGFACE network, MT-VGG. The predictions for all tasks are pooled from the same feature space.

CNN-RNN Multi-Task As shown in the experimental section, MT-VGG has the best performance; thus we construct a CNN-RNN multi-task network, based on MT-VGG. In more detail, a 2-layer GRU with 128 cells each is stacked on top of the first fc layer of MT-VGG for capturing the temporal dynamics; the output layer is on top of the GRU. We call this network MT-VGG-RNN.

CNN-RNN Multi-Modal (A/V) & Multi-Task To handle both video and audio modalities, we use a feature level fusion strategy in our developed deep learning model, that we illustrate in Figure 4. This model consists of two identical streams that extract features directly from raw input images and spectrograms, respectively. Each stream consists of a MT-VGG-RNN, described above, without the output layer. The features from the two streams are concatenated, forming a 256-dimensional feature vector that is passed through a 2-layer GRU layer with 128 units in each layer, in order to fuse the information of the audio and visual streams. The output layer follows on top of it. We call this network A/V-MT-VGG-RNN.

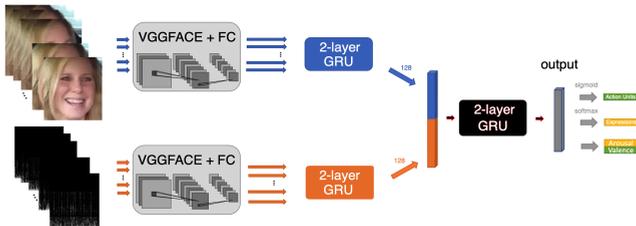


Figure 4: A/V-MT-VGG-RNN: the Multi-Modal and Multi-Task developed model

Standard Loss Functions The objective function minimized during training of the multi-task networks is the sum of the individual task losses:

$$\mathcal{L}_{CCE} = \mathbb{E}[-\log e^{p_p} + \log \sum_{i=1}^7 e^{p_i}] \quad (1)$$

$$\mathcal{L}_{BCE} = \mathbb{E}[-\sum_{i=1}^{17} (t_i \cdot \log p_i + (1 - t_i) \cdot \log (1 - p_i))] \quad (2)$$

$$\mathcal{L}_{CCC} = 1 - 0.5 \cdot (\rho_a + \rho_v), \text{ with } \rho_{a,v} = 2s_{xy} \div [s_x^2 + s_y^2 + (\bar{x} - \bar{y})^2] \quad (3)$$

where \mathcal{L}_{CCE} is the categorical cross entropy loss, \mathcal{L}_{BCE} is the binary cross entropy loss, p_p is the prediction of positive class, p_i is the prediction of AU_i , $t_i \in \{0,1\}$ is the label of AU_i , $\rho_{a,v}$ is the Concordance Correlation Coefficient (CCC) of arousal/valence, s_x and s_y are the variances of arousal/valence labels and predicted values respectively and s_{xy} is the corresponding covariance value.

ArcFace Loss Function & Networks Next, we focus on Expr recognition and introduce a new loss function. The softmax cross-entropy loss is modified as follows:

$$\mathcal{L} = \frac{-1}{N} \sum_{i=1}^N \log \frac{e^{W_{y_i}^T x_i}}{\sum_{j=1}^7 e^{W_j^T x_i}} = \frac{-1}{N} \sum_{i=1}^N \log \frac{e^{\|W_{y_i}\| \|x_i\| \cos \theta_{y_i}}}{\sum_{j=1}^7 e^{\|W_j\| \|x_i\| \cos \theta_j}} \stackrel{\|x_i\|=s, \|W_j\|=1}{=} \frac{-1}{N} \sum_{i=1}^N \log \frac{e^{s \cdot \cos \theta_{y_i}}}{\sum_{j=1}^7 e^{s \cdot \cos \theta_j}} \quad (4)$$

where the embedding feature $x_i \in \mathcal{R}^d$ denotes the deep feature of the i -th sample belonging to the y_i -th class, $W_j \in \mathcal{R}^d$ denotes the j -th column of the weight $W \in \mathcal{R}^{d \times 7}$, N is the batch size, θ_j is the angle between weight W_j and feature x_i , $\|W_j\|$ is fixed to 1 by l_2 normalization, $\|x_i\|$ is fixed by l_2 normalization and re-scaled to s .

From eq.4, it can be seen that the embedding features are distributed around each feature centre on the hypersphere. In our case, we adopt the ArcFace loss, where an angular margin penalty m between x_i and W_{y_i} is added to simultaneously enhance the intra-class compactness and inter-class discrepancy (eq.4: $\theta_{y_i} \rightarrow \theta_{y_i} + m$). m is equal to the geodesic distance margin penalty in the normalised hypersphere. We refer the interested reader to [8] for more details and explanation of this loss.

Next, we develop two networks to account for this loss. The first CNN architecture, called Multi-Task-ArcFace-Residual (MT-ArcRes) uses residual units and is depicted in Fig.5; ‘bn’ stands for batch normalization, the convolution layer is in the format: filter height \times filter width conv., number of output feature maps; the stride is equal to 2, everywhere; the fc layer is the embedding layer; the output layer provides the seven expression class logits ($W_j^T x_i$, $j = 1..7$). The second network is called Multi-Task-Arcface-VGG (MT-ArcVGG); the difference with MT-ArcRes is that the rectangular area in the Figure contains VGGFace’s layers.

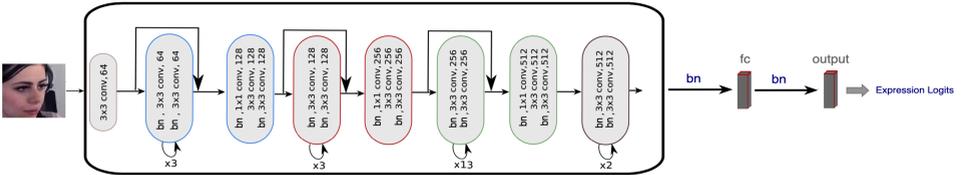


Figure 5: The MT-ArcRes network that has been trained with the ArcFace loss

4 Experimental Study

The experimental study consists of two parts. In the first, we train MT CNNs, CNN-RNNs & Multi-Modal CNN-RNNs, for VA, AU and Expr Recognition on the Aff-Wild2; then, we test the networks on 10 different databases, showing that Aff-Wild2 and the Multi-Task networks provide the best pre-trained framework for a large variety of emotion recognition settings. In the second, focusing on expression recognition, we first train ArcFace networks with Aff-Wild2 and then re-train them with each expression database; we evaluate them, achieving state-of-the-art performance.

Implementation Details & Settings Specific details about hyperparameters of the developed architectures can be found in Table 3. All experiments in this paper are

implemented in TensorFlow, on a Tesla V100 32GB GPU, using Adam optimizer (with default values) or SGD with momentum (0.9) in the ArcFace Networks' case. Additional details follow:

I) CNN Single- & Multi-Task: The networks have first been pre-trained for VA estimation on the Aff-Wild database, then the output layer is discarded and substituted by a new one for single- or multi-task, depending on the network type. Then they are trained end-to-end on Aff-Wild2.

II) CNN-RNN Multi-Task: The CNN part is initialized with the weights of the CNN MT-VGG. Then the whole architecture is trained end-to-end on Aff-Wild2.

III) CNN-RNN Multi-Modal (A/V) & Multi-Task: Training is divided in two phases: first the audio/visual streams are trained independently and then the audiovisual network is trained end-to-end. To train each stream individually, we follow the same procedure as in the CNN-RNN Multi-Task case. Once the single streams are trained, they are used for initializing the corresponding streams in the multi-stream architecture. Finally, the entire audiovisual network is trained end-to-end.

IV) ArcFace Networks: Both networks are first trained on Aff-Wild2. Then, they are re-trained end-to-end on each of the examined databases. During testing we keep the feature embedding layer, discarding the output layer. For all training images, we extract features from the embedding layer and split them in 7 clusters. Then, for each test image, we compute its distance (based on cosine similarity) from all cluster centers and assign it to the center for which this distance is minimum.

Table 3: Network Configurations: ST= Single Task, MT=Multi-Task

	ST- & MT-CNN	MT-CNN-RNN	A/V-MT-CNN-RNN	MT-ArcRes / MT-ArcVGG
learning rate	$[10^{-4}, 10^{-5}], best : 10^{-4}$	$[10^{-4}, 10^{-6}], best : 10^{-5}$	$[10^{-3}, 10^{-6}], best : 10^{-5}$	$[10^{-4}, 10^{-5}], best : 10^{-4}$
batch size/seq.length	256 / -	10 / 90	5 / 90	300 / -
parameters	dropout=0.4	dropout=0.4	dropout=0.4	dropout=0.4, $d \in \{32, 512\}, s \in \{32, 64\}, m \in \{0.1, 0.5, 1, 1.5, 2, 2.5, 3\}, best : 0.1/1$

Databases Table 4 shows the databases used in our experiments along with their properties. BP4DS and BP4D+ datasets correspond to the ones used in the FERA 2015 [10] and 2017 [12] Challenges, respectively. All databases are in-the-wild, apart from DISFA, BP4DS, BP4D+, which are spontaneous. Let us note that for the AffectNet, BP4DS and BP4D+ databases, the test set is not released; thus we report the performances on the validation set, which we use for testing.

Table 4: Properties of Databases used in our Experiments

Databases	AFEW-VA [10]	AffectNet	RAF-DB [11]	FER2013 [13]	IMFDB [14]	Emotionet	DISFA [15]	BP4DS [16]	BP4D+ [17]
Model of Affect	VA	VA, Expr	Expr	Expr	Expr	AUs	AUs	AUs	AUs
# of videos	600	-	-	-	-	-	54	1,640	5,463
# of frames	30,050	450,000	15,200	35,887	34,512	50,000	261,630	222,573	967,570

Evaluation Metrics CCC, defined in eq.3 is used for VA estimation, as it has been the evaluation criterion in all related Challenges [20, 54]. The usual F1 score is adopted for evaluation of AU detection and Expr classification. Exceptions are the RAF-DB and FER2013, in which the mean diagonal value of the confusion matrix and the accuracy metric, respectively, are the default performance measures.

Results on static databases for VA & Expr Recognition Table 5 presents the results of different CNN Single- and Multi-task (ST- and MT-) networks in a cross-database setting (networks are trained on Aff-Wild2 and tested on AffectNet, RAF-DB, FER2013 and IMFDB). The MT-VGG has the best performance for both VA and Expr recognition. In Table 5, we also compare MT-VGG’s performance with that of the state-of-the-art in each of the tested databases (the results shown are taken from the respective papers). It can be seen that MT-VGG beats the state-of-the-art in all databases, illustrating the excellent cross-performance of the generated framework. Only, in expression recognition in AffectNet, the obtained performance is lower to the state-of-the-art.

Table 5: Cross-database evaluation (models trained on Aff-Wild2 and tested on other databases) for VA and Expr on static databases: VA evaluation is shown as (CCC_V-CCC_A) ; single values correspond to expressions’ performance metrics

Databases	MT-VGG	ST-VGG	MT-SphereFace	ST-SphereFace	MT-Inc. ResNet	ST-Inc. ResNet	AlexNet [14]	VGGFACE [15]	VGG [16]
FER2013	0.76	0.73	0.72	0.72	0.74	0.71	-	-	0.75
RAF-DB	0.61	0.57	0.53	0.52	0.57	0.55	-	0.58	-
IMFDB	0.42	0.39	0.39	0.38	0.4	0.39	-	-	-
AffectNet	(0.61-0.46) 0.54	(0.51-0.42) 0.52	(0.5-0.43) 0.5	(0.5-0.4) 0.51	(0.52-0.45) 0.52	(0.5-0.42) 0.51	(0.6-0.34) 0.58	-	-

Results on video databases for VA & Expr Recognition Table 6 presents the results of CNN, CNN-RNN multi-task, single- and multi-modal networks in a cross-database setting, testing on Aff-Wild, Aff-Wild2 and AFEW-VA databases. It can be seen that the MT-VGG-RNN (trained on the visual modality) displays a better performance than the MT-VGG, for VA and Expr Recognition, in all databases. Moreover, MT-VGG-RNN performs best for valence estimation when trained with the visual modality, whereas performs best for arousal when trained with the audio modality. This is because audio tends to have thematic constancy. Consider, for example, two fight sequences in a movie, one being a flashy fight scene and the other a one-sided fight with a person being injured. In both cases, arousal can be high due to loud and pronounced music, but valence will be positive in the former and negative in the latter sequence. It can also be seen that the A/V-MT-VGG-RNN outperforms the MT-VGG-RNN, illustrating that the A/V combination improves network performance, in valence, arousal and expression estimation. Table 6 compares the performance of the proposed networks with the state-of-the-art in the examined databases. It is evident that MT-VGG and MT-VGG-RNN outperform the respective state-of-the-art.

Table 6: Cross-database evaluation for VA and Expr on video databases: VA evaluation is shown as (CCC_V-CCC_A) ; single values correspond to F1 score

Databases	MT-VGG	MT-VGG-RNN visual modality	MT-VGG-RNN audio modality	A/V-MT-VGG-RNN	best CNN [14, 15]	AffWildNet [14, 15]
Aff-Wild	(0.56-0.35)	(0.60-0.45)	(0.51-0.47)	(0.62-0.49)	(0.51-0.33)	(0.57-0.43)
Aff-Wild2	(0.38-0.3) 0.4	(0.40-0.33) 0.43	(0.34-0.36) 0.43	(0.42-0.38) 0.46	(0.33-0.25) -	(0.35-0.28) -
AFEW-VA	(0.58-0.53)	(0.6-0.6)	-	-	(0.49-0.52)	(0.52-0.56)

Results for AU Detection Table 7 compares the performance between MT-VGG and state-of-the-art networks in a cross-database setting among Emotionet, DISFA, BP4DS and BP4D+; all reported results are for the common AUs between the testing

database and Aff-Wild2. It is clear that MT-VGG outperforms the winner [9] of Emotionet 2017 Challenge, the baseline [40] and the winner [49] of FERA 2015, the baseline [42] of FERA 2017 and the fine-tuned VGG (FVGG) and R-TI method of [23]. Apart from Emotionet, in all other cases, there is a boost in performance. MT-VGG displays a slightly worse performance than the winner [58] of FERA 2017.

Table 7: Cross-database evaluation for AU Detection: evaluation metric is F1 score

Databases	MT-VGG	MT-VGG-RNN	[9]	[40]	[49]	[58]	[23]	FVGG [42]	R-TI [23]
Aff-Wild2	0.42	0.44	-	-	-	-	-	-	-
Emotionet	0.52	-	0.51	-	-	-	-	-	-
DISFA	0.61	-	-	-	-	-	-	0.52	0.60
BP4DS	0.66	-	-	0.54	0.53	-	-	-	-
BP4D+	0.49	-	-	-	-	0.51	0.34	-	-

From the above Tables, it is evident that Aff-Wild2 constitutes a very rich database for deep network training and further testing on very different and diverse emotion databases; the presented cross-database results validate our network developments.

Results with ArcFace Loss for Expr Recognition Table 8 presents a performance comparison between: i) MT-VGG (trained on Aff-Wild2), ii) a fine-tuned MT-VGG (FT-MT-VGG; pre-trained on Aff-Wild2, then re-trained on each of the examined databases), iii) the two networks trained with the ArcFace loss (MT-ArcRes, MT-ArcVGG) on Aff-Wild2 and re-trained on each of the examined databases, iv) the state-of-the-art in these databases (whose results are taken from the respective papers). The FT-MT-VGG outperforms the state-of-the-art in all databases, apart from RAF-DB, where DLP-CNN performs better, however, trained with a locality preserving loss function. Table 8 also compares this network’s performance to the performance of MT-ArcRes and MT-ArcVGG networks, trained with the ArcFace loss function. These networks outperform all others, including DLP-CNN.

Table 8: Retrained multi-task networks with ArcFace loss, for expression recognition

Databases	MT-ArcRes	MT-ArcVGG	FT-MT-VGG	MT-VGG	AlexNet [49]	DLP-CNN [24]	VGG [23]
AffectNet	0.63	0.62	0.59	0.54	0.58	-	-
RAF-DB	0.75	0.76	0.71	0.61	-	0.74	-
IMFDB	0.55	0.56	0.51	0.42	-	-	-
FER2013	0.8	0.79	0.78	0.76	-	-	0.75

5 Conclusions

In this paper, we present the first, largest, in-the-wild, A/V database, called Aff-Wild2, that is annotated for VA, AUs and Exprs. We build and train multi-task and multi-modal CNNs and CNN-RNNs on Aff-Wild2 and test their performances on 10 databases, beating the state-of-the-art. We further train two new networks on Aff-Wild2, adopting the ArcFace loss function, and then re-train them on a variety of expression databases; the results improve the existing state-of-the-art.

Acknowledgements We would like to thank Viktoriia Sharmanska for our fruitful conversations during preparation of this work. The work of S. Zafeiriou has been partially funded by the EPSRC Fellowship Deform (EP/S010203/1). The work of Dimitrios Kollias was funded by a Teaching Fellowship of Imperial College London.

References

- [1] Niki Aifanti, Christos Papachristou, and Anastasios Delopoulos. The mug facial expression database. In *11th International Workshop on Image Analysis for Multimedia Interactive Services WIAMIS 10*, pages 1–4. IEEE, 2010.
- [2] C.F. Benitez-Quiroz, R. Srinivasan, and A.M. Martinez. Emotionet: An accurate, real-time algorithm for the automatic annotation of a million facial expressions in the wild. In *Proceedings of IEEE International Conference on Computer Vision & Pattern Recognition (CVPR'16)*, Las Vegas, NV, USA, June 2016.
- [3] Sanjay Bilakhia, Stavros Petridis, Anton Nijholt, and Maja Pantic. The mahnob mimicry database: A database of naturalistic human interactions. *Pattern recognition letters*, 66:52–61, 2015.
- [4] Sumit Chopra, Raia Hadsell, Yann LeCun, et al. Learning a similarity metric discriminatively, with application to face verification. In *CVPR (1)*, pages 539–546, 2005.
- [5] Roddy Cowie and Randolph R Cornelius. Describing the emotional states that are expressed in speech. *Speech communication*, 40(1):5–32, 2003.
- [6] Roddy Cowie, Ellen Douglas-Cowie, Susie Savvidou*, Edelle McMahon, Martin Sawey, and Marc Schröder. ‘feeltrace’: An instrument for recording perceived emotion in real time. In *ISCA tutorial and research workshop (ITRW) on speech and emotion*, 2000.
- [7] Tim Dalgleish and Mick Power. *Handbook of cognition and emotion*. John Wiley & Sons, 2000.
- [8] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. *arXiv preprint arXiv:1801.07698*, 2018.
- [9] Wan Ding, Dong-Yan Huang, Zhuo Chen, Xinguo Yu, and Weisi Lin. Facial action recognition using very deep networks for highly imbalanced class distribution. In *2017 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pages 1368–1372. IEEE, 2017.
- [10] Shichuan Du, Yong Tao, and Aleix M Martinez. Compound facial expressions of emotion. *Proceedings of the National Academy of Sciences*, 111(15):E1454–E1462, 2014.
- [11] Paul Ekman. Facial action coding system (facs). *A human face*, 2002.
- [12] Mariana-Iuliana Georgescu, Radu Tudor Ionescu, and Marius Popescu. Local learning with deep and handcrafted features for facial expression recognition. *IEEE Access*, 7:64827–64836, 2019.
- [13] Ian J Goodfellow, Dumitru Erhan, Pierre Luc Carrier, Aaron Courville, Mehdi Mirza, Ben Hamner, Will Cukierski, Yichuan Tang, David Thaler, Dong-Hyun Lee, et al. Challenges in representation learning: A report on three machine

- learning contests. In *International Conference on Neural Information Processing*, pages 117–124. Springer, 2013.
- [14] Ralph Gross, Iain Matthews, Jeffrey Cohn, Takeo Kanade, and Simon Baker. Multi-pie. *Image and Vision Computing*, 28(5):807–813, 2010.
- [15] Elad Hoffer and Nir Ailon. Deep metric learning using triplet network. In *International Workshop on Similarity-Based Pattern Recognition*, pages 84–92. Springer, 2015.
- [16] Dimitrios Kollias and Stefanos Zafeiriou. Aff-wild2: Extending the aff-wild database for affect recognition. *arXiv preprint arXiv:1811.07770*, 2018.
- [17] Dimitrios Kollias and Stefanos Zafeiriou. A multi-task learning & generation framework: Valence-arousal, action units & primary expressions. *arXiv preprint arXiv:1811.07771*, 2018.
- [18] Dimitrios Kollias and Stefanos Zafeiriou. Training deep neural networks with different datasets in-the-wild: The emotion recognition paradigm. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2018.
- [19] Dimitrios Kollias, Mihalis A Nicolaou, Irene Kotsia, Guoying Zhao, and Stefanos Zafeiriou. Recognition of affect in the wild using deep neural networks. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*, pages 1972–1979. IEEE, 2017.
- [20] Dimitrios Kollias, Panagiotis Tzirakis, Mihalis A Nicolaou, Athanasios Papaioannou, Guoying Zhao, Björn Schuller, Irene Kotsia, and Stefanos Zafeiriou. Deep affect prediction in-the-wild: Aff-wild database and challenge, deep architectures, and beyond. *International Journal of Computer Vision*, 127(6-7):907–929, 2019.
- [21] Jean Kossaifi, Georgios Tzimiropoulos, Sinisa Todorovic, and Maja Pantic. Afew-va database for valence and arousal estimation in-the-wild. *Image and Vision Computing*, 2017.
- [22] Shan Li, Weihong Deng, and JunPing Du. Reliable crowdsourcing and deep locality-preserving learning for expression recognition in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2852–2861, 2017.
- [23] Wei Li, Farnaz Abtahi, and Zhigang Zhu. Action unit detection with region adaptation, multi-labeling learning and optimal temporal fusing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1841–1850, 2017.
- [24] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphereface: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 212–220, 2017.

- [25] Patrick Lucey, Jeffrey F Cohn, Takeo Kanade, Jason Saragih, Zara Ambadar, and Iain Matthews. The extended cohn-kanade dataset (ck+): A complete dataset for action unit and emotion-specified expression. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2010 IEEE Computer Society Conference on*, pages 94–101. IEEE, 2010.
- [26] Michael J Lyons, Shigeru Akamatsu, Miyuki Kamachi, Jiro Gyoba, and Julien Budynek. The japanese female facial expression (jaffe) database. In *Proceedings of third international conference on automatic face and gesture recognition*, pages 14–16, 1998.
- [27] S Mohammad Mavadati, Mohammad H Mahoor, Kevin Bartlett, Philip Trinh, and Jeffrey F Cohn. Disfa: A spontaneous facial action intensity database. *Affective Computing, IEEE Transactions on*, 4(2):151–160, 2013.
- [28] Gary McKeown, Michel Valstar, Roddy Cowie, Maja Pantic, and Marc Schroder. The semaine database: Annotated multimodal records of emotionally colored conversations between a person and a limited agent. *IEEE Transactions on Affective Computing*, 3(1):5–17, 2011.
- [29] Ali Mollahosseini, Behzad Hasani, and Mohammad H Mahoor. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *arXiv preprint arXiv:1708.03985*, 2017.
- [30] Mahyar Najibi, Pouya Samangouei, Rama Chellappa, and Larry Davis. SSH: Single stage headless face detector. In *The IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [31] Maja Pantic, Michel Valstar, Ron Rademaker, and Ludo Maat. Web-based database for facial expression analysis. In *Multimedia and Expo, 2005. ICME 2005. IEEE International Conference on*, pages 5–pp. IEEE, 2005.
- [32] Omkar M Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. In *British Machine Vision Conference (BMVC)*, 2015.
- [33] Fabien Ringeval, Andreas Sonderegger, Jens Sauer, and Denis Lalanne. Introducing the recola multimodal corpus of remote collaborative and affective interactions. In *Automatic Face and Gesture Recognition (FG), 2013 10th IEEE International Conference and Workshops on*, pages 1–8. IEEE, 2013.
- [34] Fabien Ringeval, Björn Schuller, Michel Valstar, Jonathan Gratch, Roddy Cowie, Stefan Scherer, Sharon Mozgai, Nicholas Cummins, Maximilian Schmi, and Maja Pantic. Avec 2017–real-life depression, and affect recognition workshop and challenge. 2017.
- [35] James A Russell. Evidence of convergent validity on the dimensions of affect. *Journal of personality and social psychology*, 36(10):1152, 1978.
- [36] Shankar Setty, Moula Husain, Parisa Beham, Jyothi Gudavalli, Menaka Kandasamy, Radhesyam Vaddi, Vidyagouri Hemadri, JC Karure, Raja Raju, B Rajan, et al. Indian movie face database: a benchmark for face recognition under wide variations. In *2013 Fourth National Conference on Computer Vision, Pattern Recognition, Image Processing and Graphics (NCVPRIPG)*, pages 1–5. IEEE, 2013.

- [37] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *AAAI*, volume 4, page 12, 2017.
- [38] Chuangao Tang, Wenming Zheng, Jingwei Yan, Qiang Li, Yang Li, Tong Zhang, and Zhen Cui. View-independent facial action unit detection. In *2017 12th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2017)*, pages 878–882. IEEE, 2017.
- [39] Ying-li Tian, Takeo Kanade, and Jeffrey F Cohn. Recognizing action units for facial expression analysis. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 23(2):97–115, 2001.
- [40] Michel Valstar and Maja Pantic. 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, 2010.
- [41] Michel F Valstar, Timur Almaev, Jeffrey M Girard, Gary McKeown, Marc Mehu, Lijun Yin, Maja Pantic, and Jeffrey F Cohn. Fera 2015-second facial expression recognition and analysis challenge. In *Automatic Face and Gesture Recognition (FG), 2015 11th IEEE International Conference and Workshops on*, volume 6, pages 1–8. IEEE, 2015.
- [42] Michel F Valstar, Enrique Sánchez-Lozano, Jeffrey F Cohn, László A Jeni, Jeffrey M Girard, Zheng Zhang, Lijun Yin, and Maja Pantic. Fera 2017-addressing head pose in the third facial expression recognition and analysis challenge. In *2017 12th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2017)*, pages 839–847. IEEE, 2017.
- [43] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5265–5274, 2018.
- [44] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *European conference on computer vision*, pages 499–515. Springer, 2016.
- [45] CM Whissel. The dictionary of affect in language, emotion: Theory, research and experience: vol. 4, the measurement of emotions, r. *Plutchik and H. Kellerman, Eds., New York: Academic*, 1989.
- [46] Shuo Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. Wider face: A face detection benchmark. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [47] Lijun Yin, Xiaozhou Wei, Yi Sun, Jun Wang, and Matthew J Rosato. A 3d facial expression database for facial behavior research. In *Automatic face and gesture recognition, 2006. FGR 2006. 7th international conference on*, pages 211–216. IEEE, 2006.

- [48] Lijun Yin, Xiaochen Chen, Yi Sun, Tony Worm, and Michael Reale. A high-resolution 3d dynamic facial expression database. In *Automatic Face & Gesture Recognition, 2008. FG'08. 8th IEEE International Conference On*, pages 1–6. IEEE, 2008.
- [49] Anıl Yüce, Hua Gao, and Jean-Philippe Thiran. Discriminant multi-label manifold embedding for facial action unit detection. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, volume 6, pages 1–6. IEEE, 2015.
- [50] Stefanos Zafeiriou, Dimitrios Kollias, Mihalis A Nicolaou, Athanasios Papaioannou, Guoying Zhao, and Irene Kotsia. Aff-wild: Valence and arousal‘in-the-wild’challenge. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 34–41, 2017.
- [51] Xing Zhang, Lijun Yin, Jeffrey F Cohn, Shaun Canavan, Michael Reale, Andy Horowitz, Peng Liu, and Jeffrey M Girard. Bp4d-spontaneous: a high-resolution spontaneous 3d dynamic facial expression database. *Image and Vision Computing*, 32(10):692–706, 2014.
- [52] Zheng Zhang, Jeff M Girard, Yue Wu, Xing Zhang, Peng Liu, Umur Ciftci, Shaun Canavan, Michael Reale, Andy Horowitz, Huiyuan Yang, et al. Multimodal spontaneous emotion corpus for human behavior analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3438–3446, 2016.