

AVEC 2019 Workshop and Challenge: State-of-Mind, Detecting Depression with AI, and Cross-Cultural Affect Recognition

Fabien Ringeval
Université Grenoble Alpes, CNRS
Grenoble, France

Björn Schuller*
University of Augsburg
Augsburg, Germany

Michel Valstar
University of Nottingham
Nottingham, UK

Nicholas Cummins
University of Augsburg
Augsburg, Germany

Roddy Cowie
Queen's University Belfast
Belfast, UK

Leili Tavabi
University of Southern California
Los Angeles, USA

Maximilian Schmitt
University of Augsburg
Augsburg, Germany

Sina Alisamir
Université Grenoble Alpes, CNRS
Grenoble, France

Shahin Amiriparian
University of Augsburg
Augsburg, Germany

Eva-Maria Messner
University of Ulm
Ulm, Germany

Siyang Song
University of Nottingham
Nottingham, UK

Shuo Liu
University of Augsburg
Augsburg, Germany

Ziping Zhao
Tianjin Normal University
Tianjin, China

Adria Mallol-Ragolta
University of Augsburg
Augsburg, Germany

Zhao Ren
University of Augsburg
Augsburg, Germany

Mohammad Soleymani
University of Southern California
Los Angeles, USA

Maja Pantic[†]
Imperial College London
London, UK

ABSTRACT

The Audio/Visual Emotion Challenge and Workshop (AVEC 2019) “State-of-Mind, Detecting Depression with AI, and Cross-cultural Affect Recognition” is the ninth competition event aimed at the comparison of multimedia processing and machine learning methods for automatic audiovisual health and emotion analysis, with all participants competing strictly under the same conditions. The goal of the Challenge is to provide a common benchmark test set for multimodal information processing and to bring together the health and emotion recognition communities, as well as the audiovisual processing communities, to compare the relative merits of various approaches to health and emotion recognition from real-life data. This paper presents the major novelties introduced this year, the challenge guidelines, the data used, and the performance of the

baseline systems on the three proposed tasks: state-of-mind recognition, depression assessment with AI, and cross-cultural affect sensing, respectively.

CCS CONCEPTS

• **General and reference** → **Performance**;

KEYWORDS

Affective Computing; State-of-Mind; Cross-Cultural Emotion

ACM Reference Format:

Fabien Ringeval, Björn Schuller, Michel Valstar, Nicholas Cummins, Roddy Cowie, Leili Tavabi, Maximilian Schmitt, Sina Alisamir, Shahin Amiriparian, Eva-Maria Messner, Siyang Song, Shuo Liu, Ziping Zhao, Adria Mallol-Ragolta, Zhao Ren, Mohammad Soleymani, and Maja Pantic. 2019. AVEC 2019 Workshop and Challenge: State-of-Mind, Detecting Depression with AI, and Cross-Cultural Affect Recognition. In *9th International Audio/Visual Emotion Challenge and Workshop (AVEC '19)*, October 21, 2019, Nice, France. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3347320.3357688>

*The author is further affiliated with Imperial College London, London, UK.

[†]The author is further affiliated with University of Twente, Twente, The Netherlands.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

AVEC '19, October 21, 2019, Nice, France

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6913-8/19/10...\$15.00

<https://doi.org/10.1145/3347320.3357688>

1 INTRODUCTION

The Audio/Visual Emotion Challenge and Workshop (AVEC 2019) is the ninth competition aimed at comparison of multimedia processing and machine learning methods for automatic audio, visual, and audiovisual health and emotion sensing, with all participants competing strictly under the same conditions [48–50, 62, 63, 72, 74, 75].

One of the goals of the AVEC series is to bring together multiple communities from different disciplines, in particular, the audiovisual multimedia communities and those in the psychological and social sciences who study expressive behaviour. Another objective is to advance health and emotion recognition systems by providing a common benchmark test set for multimodal information processing, in order to compare the relative merits of the approaches to automatic health and emotion analysis under well-defined conditions, i. e. , with large volumes of un-segmented, non-prototypical and non-preselected data of wholly naturalistic behaviour. This is precisely the type of data that the new generation of affect-oriented multimedia and human-machine/human-robot communication interfaces have to face in the real world.

Major novelties are introduced for the AVEC 2019 with three separated Sub-challenges focusing on health and emotion analysis: (i) State-of-Mind Sub-challenge (SoMS), (ii) Detecting Depression with AI Sub-challenge (DDS), and (iii) Cross-cultural Emotion Sub-challenge (CES). In the following, we describe the novelties introduced in the Challenge and the guidelines for participating.

The State-of-Mind Sub-challenge (SoMS) is a new task focusing on the continuous adaptation of human state-of-mind (SOM), which is pivotal for mental functioning and behaviour regulation [25]. SOM is constantly shifting due to internal and external stimuli, and frequent use of either adaptive or maladaptive SOM influences our mental health. One key aspect of the human experience is the way emotion features in our SOM [64, 68]. In the SoMS, self-reported mood (10-point Likert scale) after the narrative of personal stories (two positive and two negative), has to be predicted automatically from the audiovisual recordings of those stories; USoM corpus [46].

The Detecting Depression with AI Sub-challenge (DDS) is a major extension of the AVEC 2016 DSC [73], where the level of depression severity (PHQ-8 questionnaire) was assessed from audiovisual recordings of patients interacting with a virtual agent conducting a clinical interview and driven by a human as a Wizard-of-Oz (WoZ); DAIC-WOZ corpus [23]. The DAIC data set contains new recordings of the same population with the virtual agent being, this time, wholly driven by AI, i. e. , without any human intervention. Those new recordings are used as a test partition for the DDS, and will help to understand how the absence of a human conducting the virtual agent impacts on automatic depression severity assessment.

The Cross-cultural Emotion Sub-challenge (CES) is a large extension of the AVEC 2018 CES [47], where dimensions of emotion were inferred from audiovisual recordings collected “*in-the-wild*”, i. e. , with standard webcams and at home/work place. A cross-cultural setup was further exploited for inferring emotion: knowledge of German culture was leveraged to infer emotion on the Hungarian culture, using the SEWA corpus [31]. This dataset now includes data collected from new participants with Chinese culture, which is used as a test set for the CES, whose aim is, therefore, to investigate how emotion knowledge of Western European cultures (German, Hungarian) can be transferred to the Chinese culture.

All Sub-challenges allow contributors to find their own features to use with their own machine learning algorithm. In addition, standard feature sets are provided for audio and video data (cf. Section 4),

along with scripts available in a public repository¹, which participants are free to use for reproducing both the baseline features and recognition systems (cf. Section 5). The labels of the test partition remain unknown to the participants, and participants have to stick to the definition of training, development, and test partition. They may freely report on results obtained on the development partition, but are limited to five trials per Sub-challenge in submitting their results on the test partition.

Ranking of the labels relies on the *Concordance Correlation Coefficient (CCC)* [36] for all Sub-challenges; the Root Mean Squared Error (*RMSE*) is additionally reported. Whereas many other metrics of performance could be exploited for ranking the contributions, such as the Spearman’s *CC*, or the coefficient of determination (r^2), we believe that the index of reproducibility *CCC* is the most suitable metric to use, as it is not biased by changes in scale and location, and elegantly includes information on both precision and accuracy in a single statistical measure [36]. Moreover, its theoretical definition and properties are well rooted in the literature [41], and it can be easily exploited as a loss function for training neural networks [77].

To be eligible to participate in the Challenge, every entry has to be accompanied by a paper submitted to the AVEC 2019 Data Challenge and Workshop, describing the results and the methods that created them. These papers undergo peer-review by the technical program committee. Only contributions with a relevant accepted paper and at least a submission of test results are eligible for participation. The organisers do not participate in the Challenge themselves, but re-evaluate the findings of the best performing system of each Sub-challenge.

The remainder of this paper is organised as follows. We summarise relevant related work in Section 2, introduce the Challenge corpora in Section 3, the common audiovisual baseline feature sets in Section 4, and the developed baseline recognition systems with the obtained results in Section 5, before concluding in Section 6.

2 RELATED WORK

This section is a summary of the current state-of-the-art in the automatic analysis of affect with a focus on: (i) human state-of-mind, (ii) depression assessment in the context of AI-driven virtual agents, and (iii) dimensional analysis in cross-cultural paradigms.

2.1 State-of-Mind

The concept of a human SOM describes the phenomenon that our consciousness and emotions are constantly fluctuating over time; this is due to internal and external biological, psychological, and social demands [25, 46]. Our emotions provide valuable information that influences our basic human processes in a bidirectional manner [67, 68]. Such processes include attention, perception, cognition, memory retrieval, memory storage, and behaviour regulation. In fact, depending on our actual SOM, some emotions, cognitions, and behaviours are more likely to occur, while others may be suppressed. This effect is the underlying principle of mood congruence [52, 64].

Despite the major impact of SOM on health and social functioning, the quantification of current emotional states, with therapy contexts, has its pitfalls. The simplest of these is that it relies heavily on self-reports of emotional states, which are inherently

¹<https://github.com/AudioVisualEmotionChallenge/AVEC2019>

biased [79]. It is indeed not always sound to assume that people who give the same scores on measurement scales are actually in the same SOM [38]. Moreover, even within a person, the current rating of the emotional state is rooted in previous experiences, known as the adaptation level, and therefore is not really accurate in an absolute way [53].

Approaches like Russell's avoid having to limit the quantification to a given language [52]: on his theory of core affect, every instance of emotion can be quantified on the orthogonal axes *arousal* (from sleepy to hyper-aroused) and *valence* (with the poles negative and positive). However, raters' ability to quantify reliably is still doubtful. One approach to overcoming this is to treat emotional state values as ordinal variables [79]. Another is to complement self-ratings with expert ratings or physiological recordings. Each of these methods has its limitations; the mismatch between different emotion assessments is still very much a matter of scientific discourse [65, 66].

Despite the given limitations of the scientific assessment of emotional states, humans constantly monitor their own and others' emotions and organise themselves within social systems [12, 55]. Given the need for humans to socially interact and the increased occurrence of human-machine-interactions, the development of a real-time SOM data-driven recognition system has the potential to enhance user experience, user satisfaction, and subsequently to foster user adherence [5, 45, 46]. Such a system could assist society in various ways; i) decreasing bias in the monitoring of SOM; ii) collecting more objective data to aid the diagnosis of affective disorders; iii) delivering tailored interventions to facilitate treatment of disease; iv) reducing the time spent in the evaluation of treatment outcome, and in e-treatment by presenting SOM related content, easing burdens on both patient and provider [45, 46, 61, 70].

2.2 Depression Detection with AI

Depression, particularly major depressive disorder (MDD), is a common mental health problem, with negative impacts on the way one thinks, feels, and acts [3]. It can lead to a variety of emotional and physical problems and affect many aspects of both working and personal life. The World Health Organisation (WHO) declared depression as the leading cause of ill health and disability worldwide in 2015: more than 300 million people live with it [40]. Given the high prevalence of depression and its suicide risk, finding new methods for diagnosis and treatment becomes more and more critical.

There is growing interest in using automatic human behaviour analysis for computer-aided depression diagnosis based on behavioural cues such as facial expressions and speech prosody, because of convincing evidences that depression and related mental health disorders are associated with changes in patterns of behaviour [8, 11, 28, 58, 78]. Facial activity, gesturing, head movements and expressivity are among behavioural signals that are strongly correlated with depression.

Early paralinguistic investigations into depressed speech found that patients consistently demonstrated prosodic speech abnormalities such as reduced pitch, reduced pitch range, slower speaking rate, and higher articulation errors [11]. Facial expression and head gestures that can be tracked by computer vision are also good predictors of depression; e. g., a more downward angle of the gaze, less

intense smiles, and shorter average duration of smiles have been reported as the most salient facial cues of depression [57]. Further, body expressions, gestures, head movements, and linguistic cues have also been reported to provide relevant cues for depression detection [1, 39, 43, 44].

Taking all those evidences together, it has been proposed to integrate affective computing technology into a computer agent that interviews people and identifies verbal and nonverbal indicators of mental illnesses [14]. Data collected with subjects suffering from post-traumatic stress disorder showed that the automatic evaluation of their level of depression severity (PHQ-8 questionnaire) can achieve a *RMSE* less than 5 when the agent is driven by a human acting as a WoZ [22]; PHQ-8's range $\in [0, 24]$ and cutpoints are defined at [5, 10, 15, 20] for mild, moderate, moderately severe, and severe depression, respectively. Those results need to be investigated further, with the agent being wholly driven by AI, as the wizard might drive the virtual agent to a situation that eases the observation of patterns associated with depression, or the autonomous agent might have issues in conducting the interview appropriately.

2.3 Cross-cultural Emotion Recognition

Cross-cultural emotion recognition has long been highlighted as an open research question within the affective computing community [15, 17, 18, 42], and was introduced as an AVEC Sub-challenge in 2018 [47]. Whereas the AVEC 2018 CES focused on detecting *arousal*, *valence*, and *liking* from Hungarian speakers using only German speakers for training and development of the models [47], in this year's AVEC CES the test cohort is Chinese speakers with speakers from the two cultures mentioned earlier being available for training, development, and additional testing.

A common belief in facial expression recognition is that emotional expressions have a large degree of universality across cultures [10, 16]. This statement was on the whole supported by both baseline results and works submitted to the AVEC 2018 CES, with either vision-only or multimodal systems achieving higher cross-culture accuracies than speech-only approaches [27, 49, 76, 82]. These results were insightful, as previously, there were only a few works in the affective computing literature which supported this claim [9, 15].

Interestingly, approaches in the AVEC 2018 CES did not employ approaches such as transfer learning [80, 81] or domain adaptation techniques [29, 54] typically seen in cross-cultural testing. In [76], the authors proposed a model based on emotional salient detection to identify emotion markers invariant to socio-cultural context. The other two entrants employed data driven approaches based on long short-term memory recurrent neural networks (LSTM-RNN) [27, 82]. Matching with similar results in the literature [21, 56], all entrants in the AVEC 2018 CES observed a drop in system performance when testing on the Hungarian data [27, 76, 82].

3 CHALLENGE CORPORA

The AVEC 2019 Challenge relies on three corpora: (i) the USoM corpus [46] for the SoMS, (ii) the Extended-DAIC corpus [23] for the DDS, and (iii) the SEWA dataset [31] for the CES. We provide below a short overview of each dataset and refer the reader to the original work for a more complete description.

3.1 Ulm State-of-Mind Corpus

The Ulm state of mind database was collected to assess the association between personal story telling and current SOM, operationalised by affective state according to Russel’s theory [46, 52, 61]. Parts of this dataset have been released for the Interspeech 2018 Computational Paralinguistics (ComParE) challenge [61].

Participants of the USoM corpus were instructed to first tell two negative personal narratives $NN_{1,2}$ and subsequently two positive personal narratives $PN_{1,2}$, each for five minutes in front of a camera. They were also asked to rate their current affect (CA) on a 10-point likert scale for the dimensions *arousal* and *valence* before and after telling each narrative, resulting in the following protocol: (t_0), $CA_0, NN_1, (t_1), CA_1, NN_2, (t_2), CA_2, PN_1, (t_3), CA_3, PN_2, (t_4), CA_4$. For the purpose of the Challenge, the USoM dataset was partitioned into training, development, and test sets while preserving the overall speaker diversity – in terms of age, gender distribution, and core affect evaluations – within the partitions. Table 1 shows the number of subjects and duration for each partition.

As the interest of the SoMS is on the change in mood, rather than just its static observation, the initial self-reports made before the storytelling are included in the data package given to participants for all partitions, including the test set. Exploiting such contextual information in an automatic system predicting the level of mood is a realistic scenario in the real-world, because a therapist usually asks a person’s baseline emotion at the start of a session. It is thus essential to provide machine learning algorithms with the same prior information as a therapist might have.

3.2 Distress Analysis Interview Corpus

The Extended Distress Analysis Interview Corpus (E-DAIC) [14] is an extended version of WOZ-DAIC [23] that contains semi-clinical interviews designed to support the diagnosis of psychological distress conditions such as anxiety, depression, and post-traumatic stress disorder. These interviews were collected as part of a large effort to create a computer agent that interviews people and identifies verbal and nonverbal indicators of mental illnesses [23].

Data collected include audio and video recordings, automatically transcribed text using Google Cloud’s speech recognition service, and extensive questionnaire responses. The interviews are conducted by an animated virtual interviewer called Ellie. In the WoZ interviews, the virtual agent is controlled by a human interviewer (wizard) in another room, whereas in the AI interviews, the agent acts in a fully autonomous way using different automated perception and behaviour generation modules.

For the purpose of the Challenge, the E-DAIC dataset was partitioned into training, development, and test sets while preserving the overall speaker diversity – in terms of age, gender distribution, and the eight-item Patient Health Questionnaire (PHQ-8) scores – within the partitions. Whereas the training and development sets include a mix of WoZ and AI scenarios, the test set is solely constituted from the data collected by the autonomous AI. Details regarding the speaker distribution over the partitions are given in Table 2.

Table 1: Number of subjects and duration of the storytellings contained in the USoM database [46].

Partition	# Subjects	Duration [h:min:s]
Training	45	13:49:38
Development	33	10:46:57
Test	33	9:46:14
All	111	34:22:49

Table 2: Number of subjects and duration of the interviews included in the Extended-DAIC database [23].

Partition	# Subjects	Duration [h:min:s]
Training	163	43:30:20
Development	56	14:47:31
Test	56	14:52:42
All	275	73:10:33

Table 3: Number of subjects and duration of the video chats contained in the SEWA database [31].

Culture	Partition	# Subjects	Duration [h:min:s]
German	Training	34	1:33:12
German	Devel.	14	0:37:46
German	Test	16	0:46:38
Hungarian	Training	34	1:08:24
Hungarian	Devel.	14	0:28:42
Hungarian	Test	18	0:36:06
Chinese	Test	70	3:17:52
All		200	8:28:40

3.3 Cross-cultural Emotion Database (SEWA)

The SEWA database consists of audiovisual recordings of spontaneous behaviour of participants captured using an *in-the-wild* recording paradigm [31]. Pairs of friends or relatives from German, Hungarian, and Chinese cultures were recorded through a dedicated video chat platform which utilised participants’ own – standard – web-cameras and microphones. After watching a set of commercials, pairs of participants were given the task of discussing the last advert watched (a video clip advertising a water tap) for up to three minutes. The aim of this discussion was to elicit further reactions and opinions about the advert and the product advertised.

The video chats of the three cultures have been annotated w. r. t. the emotional dimensions *arousal* and *valence*, and a third dimension describing *liking* (or sentiment), independently by several native speakers; German and Chinese: six annotators, Hungarian: five annotators. The annotation contours (traces) are combined into a single gold-standard using the same *evaluator weighted estimator* (EWE)-based approach that was used in the last two editions of AVEC [47, 51]. Table 3 shows the number of subjects and the duration of the recordings for each partition.

4 BASELINE FEATURES

Emotion recognition from audiovisual signals usually relies on feature sets whose extraction is based on expertise gained over several decades of research in the domains of speech processing, e. g., Mel Frequency Cepstral Coefficients (MFCCs), and vision computing, e. g., Facial Action Units (FAUs). However, recent advances in the field of representation learning, whose objective is to learn representations of data that are best suited for the recognition task [6], have shown that efficient representations of audiovisual signals can be learnt in the context of emotion [2, 59, 71].

Audiovisual representations can be learnt from expert-driven information extracted from the raw signals [59], or directly from the raw signals [71]. They can also be generated using adversarial networks [13], or using convolutional neural networks trained on out-of-domain data and for a different task, e. g., audio representations extracted by a model trained for object classification in images [2].

4.1 Expert-knowledge

The traditional approach in affect sensing consists in summarising low-level descriptors (LLDs) of audiovisual signals over time with a set of statistical measures computed over a fixed-duration sliding analysis window. Those descriptors usually include spectral, cepstral, prosodic, and voice quality information for the audio channel, and appearance, geometric, and FAUs information for the video channel.

As audio features, we compute the extended Geneva Minimalistic Acoustic Parameter Set (EGEMAPS) [19], which contains 88 measures covering the aforementioned acoustic dimensions, and used here as baseline. In addition, MFCCs 1-13, including their 1st- and 2nd-order derivatives (deltas and double-deltas) are computed as a set of acoustic LLDs, using the OPENSMILE² [20] toolkit. As visual features, we extract the intensities of 17 FAUs for each video frame, along with a confidence measure, using the toolkit OPENFACE³ [4]. Descriptors of pose and gaze are additionally extracted.

Audiovisual LLDs are summarised over time by computing their mean and standard-deviation using a sliding window of 4 s length, and a hop size of 1 s for the USoM and E-DAIC datasets, and 100 ms for the SEWA dataset, excepted for the EGEMAPS set, which is computed on each window.

4.2 Bags-of-Words

The technique of bags-of-words (BoW), which originates from text processing, represents the distribution of LLDs according to a dictionary learnt from them. As a front-end of the BoW, we use the MFCCs and the EGEMAPS set for the acoustic data, and the intensities of the FAUs for the video data; MFCCs and EGEMAPS LLDs are standardised (zero mean, unit variance) in an on-line approach prior to vector quantisation, while this step is not required for the FAU intensities.

To generate the BoW representations, both the acoustic and the visual features are processed and summarised over a block of a 4 s length duration, for each step of 100 ms for the SEWA dataset, and 1 s for the USoM and E-DAIC datasets. The codebook size is 100.

Instances are sampled at random to build the dictionary, and the logarithm is taken from resulting term frequencies in order to compress their range. The whole cross-modal BoW (XBOW) processing chain is executed using the open-source toolkit OPENXBOW⁴ [60].

4.3 Deep Representations

As in last year's challenge [47], we have included DEEP SPECTRUM⁵ features as a deep learning based audio baseline feature representation [2]. DEEP SPECTRUM features are inspired by deep representation learning paradigms common in image processing: spectral images of speech instances are fed into pre-trained image recognition CNNs and a set of the resulting activations are extracted as feature vectors.

For this year's challenge, we extracted DEEP SPECTRUM features from four robust pre-trained CNNs using VGG-16 [69], ALEXNET [34], DENSENET-121, and DENSENET-201 [26]; ALEXNET was used in the AVEC 2019 CES purely for consistency with the previous AVEC 2018 CES. The speech files are first transformed into mel-spectrogram images with 128 mel-frequency bands, a window width of 4 s for all challenge corpora and a hop size of 1 s for the USoM and E-DAIC datasets, and 100 ms for the SEWA dataset. Following that, the spectral-based images are forwarded through the pre-trained networks. A 4 096-dimensional feature vector is then formed from the activations of the second fully connected layer in VGG-16 and ALEXNET, and a 1 024 and a 1 920-dimensional feature vector is obtained from the activations of the last average pooling layer of the DENSENET-121 and DENSENET-201 networks, respectively.

We also provide two baseline deep visual representations. For these, we employed a VGG-16 [69] network and a ResNet-50 network [24] that are pre-trained with the Affwild dataset [30]. The pipeline starts with applying the OPENFACE toolkit [4] to detect the face region and subsequently perform face alignment. Then, we froze the weights of two pre-trained models and fed the aligned face images to both CNNs individually. To obtain the deep representations for each frame, we extract the output of the first fully-connected layer from the pre-trained VGG-16 network, and the output of the global average pooling layer from the pre-trained ResNet-50 network, respectively. As a result, a 4 096-dimensional deep feature vector from VGG and a 2 048-dimensional deep feature vector from ResNet are provided for each frame.

5 BASELINE SYSTEMS

All baseline systems rely exclusively on existing open-source machine learning toolkits to ensure the reproducibility of the results. In this section, we describe the systems developed for each Sub-challenge, and present the obtained results. For evaluation on the test set, we retained the two audio representations with the best performance, and the two video representations with the best performance, in addition to the fusion of all audiovisual representations.

²<http://audeering.com/technology/opensmile/>

³<https://github.com/TadasBaltrusaitis/OpenFace/>

⁴<https://github.com/openXBOW/openXBOW>

⁵<https://github.com/DeepSpectrum/DeepSpectrum>

Table 4: Baseline results evaluated with CCC for the AVEC 2019 SoMS; USoM data set [46]; BoAW-M/e: bags-of-audio-words with MFCCs/eGeMAPS; DS-DNet: Deep Spectrum using DenseNet-121; DS-VGG: Deep Spectrum using VGG-16; best result on the test partition is highlighted in bold.

Partition	Audio				Video				Fusion		
	MFCCs	eGeMAPS	BoAW-M	BoAW-e	DS-DNet	DS-VGG	FAUs	BoVW	ResNet	VGG	All
<i>Random sampling of training instances</i>											
Development	.282	.412	.336	.295	.280	.384	.372	.317	.261	.318	.417
Test	–	.276	–	–	–	.289	.119	–	–	.191	.278
<i>Curriculum sampling of training instances</i>											
Development	.299	.378	.334	.288	.326	.437	.419	.313	.300	.318	.464
Test	–	.294	–	–	–	.208	.151	–	–	.160	.236

5.1 State-of-Mind Sub-challenge

We use a gated recurrent unit (GRU) network with two layers, each having 64 nodes for their hidden layers, for each audiovisual representation. As a pre-processing step, all input features are normalised to have zero mean and unit variance. Dropout, at a rate of 10 %, is employed during training. The GRU is then followed by a fully connected neural network that has one hidden layer with 32 nodes, followed by a single linear layer to map to the desired output size of one. Note that a middle-fusion of the audiovisual representations is performed by concatenating their respective GRU outputs.

The model is implemented using a PYTORCH framework and is trained with an ADAM optimiser. As previous studies have shown the benefits of training a network following a curriculum [7, 37], where instances are gradually presented in increasing level of difficulty, we implemented this approach using the following strategy. First, a uniform distribution of valence labels is obtained by duplicating training instances, then, a sub-set of the training set with only the data instances with $CA \in [2-3] \cup [9-10]$, i. e., the most positive and negative storytellings, is firstly used for training, followed by a larger sub-set with data instances with $CA \in [2-4] \cup [8-10]$, each for 32 epochs. We then exploited the whole training set until *early stopping* occurs; once 60 epochs have passed, training is stopped if there is no improvement within the last 25 epochs.

Because the interest of the SoMS is in the analysis of a change in human SOM, the network is trained to model the difference between the self-reported core affect after each story and before the first story: $CA_i - CA_0$, $i = 1, 2, 3, 4$. Results are reported for each audiovisual representation, and for the two training approaches, i. e., with or without curriculum, in Table 4. Whereas the mid-fusion of all audiovisual representations provides the best result on the development set for the two learning approaches, audio descriptors achieve higher performance on the test set, with the expert-based eGeMAPS set performing best with curriculum learning.

A summary of the results obtained with either a static (CA_i) or a dynamic ($CA_i - CA_0$) view of the self-reported mood used for training or testing the system is also provided in Table 5. Interestingly, results show that the automatic inference of the self-reported mood performs much better in a ‘mixed’ scenario, i. e., training on the static view (CA_i) and evaluating on the change ($CA_i - CA_0$) or *vice-versa* training on the change and testing on the static label,

Table 5: Comparison of the approaches – training or testing on a static or dynamic measure of mood – used for the AVEC 2019 SoMS; averaged CCC results are reported; $[\mu(\sigma)]$.

Partition	Static training	Dynamic training
<i>Static evaluation</i>		
Development	.149 (.108)	.335 (.050)
Test	.037 (.063)	.219 (.068)
<i>Dynamic evaluation</i>		
Development	.368 (.150)	.102 (.066)
Test	.325 (.052)	.040 (.094)

compared to a ‘consistent’ approach with both training and testing performed on the same view, i. e., either static or dynamic.

This result might stem from emotion data being hierarchically organised. As such, each self-reported emotion is nested within a person over a period of time [32]. Because of human’s inability to assess their own emotions as an absolute value, self-reported emotion can only be interpreted as a current assessment of emotional differences in relation to the nearest past. Furthermore, there is also variance in emotion dynamics between people and not only within a person [33]. The inter-individual and intra-individual variance in emotion dynamics are strongly related to one another, but add both new information to predictions. While the variance between persons might be best captured in a scenario where machine learning is applied to raw values, the intra-individual auto-correlation of emotion, the so-called emotional inertia, is portrayed in the dynamic evaluation [35]. Therefore, training on static data and evaluating on dynamic data, such as emotional inertia, might be the *state-of-the-art* approach to characterise human SOM.

5.2 Detecting Depression Sub-challenge

For the depression detection baseline, we employ a single-layer 64-d GRU as our recurrent network with a dropout regularisation of rate 20 %, followed by a 64-d fully-connected layer to obtain a single-value regression score. To handle bias, we convert the PHQ-8 score labels to floating point numbers by downscaling with a factor of 25 prior to training. The network is trained and evaluated

Table 6: Baseline results evaluated with CCC for the AVEC 2019 DDS; RMSE is additionally reported; BoAW-M/e: bags-of-audio-words with MFCCs/eGeMAPS; DS-DNet: Deep Spectrum using DenseNet-121; DS-VGG: Deep Spectrum using VGG-16; best result on the test partition is highlighted in bold.

Partition	Audio						Video				Fusion
	MFCCs	eGeMAPS	BoAW-M	BoAW-e	DS-DNet	DS-VGG	FAUs	BoVW	ResNet	VGG	All
<i>Regression of PHQ-8 score (CCC)</i>											
Development	.198	.076	.102	.272	.165	.305	.115	.107	.269	.108	.336
Test	–	–	–	.045	–	.108	.019	–	.120	–	.111
<i>Regression of PHQ-8 score (RMSE)</i>											
Development	7.28	7.78	6.32	6.43	8.09	8.00	7.02	5.99	7.72	7.69	5.03
Test	–	–	–	8.19	–	9.33	10.0	–	8.01	–	6.37

using a CCC loss function and evaluation score, and the RMSE results are reported using the original PHQ scale. A batch size of 15 is used consistently, and the learning rate is optimised across different feature sets. In order for the data to fit on GPU memory, a maximum sequence length has been assigned for the sessions. For the MFCCs and eGeMAPS LLDs, and the high dimensional deep representations like DeepSpectrum, ResNet, and VGG, a maximum sequence length of 20 minutes is used. Additionally, for ResNet, VGG, and DEEP SPECTRUM representations frames are dropped keeping one out of two, or one out of four frames depending on the dimensionality so that the data can be loaded onto memory. Fusion of the different audiovisual representations is achieved by averaging their scores.

Baseline results of the DDS are given in Table 6. They show that, on the development set, the best CCC score from audio features was achieved with Deep spectrum (DS-VGG) features, and the model with ResNet features achieved the best result for visual features. These results indicate the power of representations learnt by deep neural networks with a large amount of data when being used in a different context to which they were initially designed, which is confirmed on the test set with the ResNet visual model achieving the best result, despite a relatively low CCC.

Fusion of the different representations achieves the best result on the development set, and the RMSE returned on the test set is slightly better than the one obtained on the DAIC-WoZ dataset with the AVEC 2017 baseline system [51]; $RMSE = 6.37$ for AVEC 2019 compared to $RMSE = 6.97$ for AVEC 2017. However, the baseline system developed for this year’s Challenge is more complex – a simple linear regression model vs GRU-RNNs for this year –, and the corresponding scores should be therefore best regarded in the light of the best results of the AVEC 2017 Depression Sub-challenge [22], which was $RMSE = 4.99$.

On the basis of the results obtained in the automatic sensing of the level of depression from interactions with the virtual agent, recognition seems more challenging when the agent is solely AI driven, than when a human is driving the agent as a WoZ. This observation opens interesting research questions for designing the agent in a way that the observation of depression cues can be maximised, e. g. , by reinforcement learning, according to the interaction style of the agent.

5.3 Cross-cultural Emotion Sub-challenge

For the baseline system of the CES, we employ a 2-layer LSTM-RNN (64 / 32 units) as a time-dependent regressor of the three targets (learnt together) for each representation of the audiovisual signals, and SVMs – LIBLINEAR with L2-L2 dual form of the objective function – for the late fusion of the predictions. The model is implemented using the KERAS framework. The network is trained for 50 epochs with the RMSPROP optimiser using a dropout rate of 10 %, and the model providing the highest CCC on the development set of the German and Hungarian culture is used to generate the predictions for the test sets (German, Hungarian, and all clips of the Chinese culture). Even though the model has three outputs modelling each dimension, the optimum model for each dimension is selected separately. The predictions of all test sequences from each culture are concatenated prior to computing the CCC, whose opposite is used as loss function for training the networks [71, 77].

In order to perform time-continuous prediction of the emotional dimensions, audiovisual signals were processed with a sliding window of 4 s length, which is a compromise to capture enough information to be used with both static regressors, such as SVMs, and context-aware regressors, such as RNNs. We utilised frame-stacking for the SVM-based late fusion of the audiovisual representations with either past, or future context.

Baseline results of the CES are given in Table 7. They show improvements over the performance reported in the previous edition of the AVEC CES; relative improvement for German is 7.25% and 8.25% for arousal and valence, respectively, and for Hungarian, 17.3% and 13.3%, respectively. The inclusion of instances of the Hungarian culture as training and development material, in addition to those of the German culture, might explain the large increase in performance for both cultures, as only instances of the German culture were available for training and development in the AVEC 2018 CES. In addition, a more recent version of the OPENFACE toolkit [4] was exploited, which provided the best results on the test set for both arousal and valence with FAUs based features. Those results confirm the common view that facial expressions of emotion have a large degree of universality across cultures, compared to the vocal expressions, where the acoustic and prosodic dimensions already play a key role in the oral communication by serving many grammatical and pragmatic functionalities, e. g. , in tonal languages

Table 7: Baseline results evaluated with CCC for the AVEC 2019 CES; SEWA dataset [31]; DeepSpec: Deep Spectrum; best result on the test partition is highlighted in bold.

Culture	Partition	Audio				Video				Fusion	
		MFCCs	eGeMAPS	BoAW-M	BoAW-e	DS	FAUs	BoVW	ResNet	VGG	All
<i>Arousal</i>											
German	Dev.	.389	.396	.323	.434	.380	.606	.556	.475	.561	.629
German	Test	–	.293	–	.276	–	.562	–	–	.505	.517
Hungarian	Dev.	.236	.305	.237	.291	.156	.425	.321	.460	.367	.583
Hungarian	Test	–	.272	–	.250	–	.527	–	–	.396	.525
Ger. + Hun.	Dev.	.326	.371	.298	.398	.312	.531	.467	.473	.493	.614
Chinese	Test	–	.100	–	.107	–	.355	–	–	.297	.238
<i>Valence</i>											
German	Dev.	.344	.405	.190	.455	.317	.639	.594	.552	.595	.684
German	Test	–	.309	–	.325	–	.627	–	–	.646	.622
Hungarian	Dev.	.017	.073	.042	.135	.084	.463	.421	.373	.363	.508
Hungarian	Test	–	.166	–	.151	.173	.459	–	–	.548	.397
Ger. + Hun.	Dev.	.187	.286	.134	.352	.233	.565	.523	.487	.505	.615
Chinese	Test	–	.267	–	.281	–	.468	–	–	.398	.423
<i>Liking</i>											
German	Dev.	.159	.136	.140	.003	.164	.056	.073	.057	.244	.048
German	Test	–	.012	–	.074	–	-.042	–	–	-.052	-.019
Hungarian	Dev.	.115	.192	-.027	.253	.121	.104	.041	.028	.028	.260
Hungarian	Test	–	.051	–	.089	–	-.062	–	–	-.069	-.22
Ger. + Hun.	Dev.	.144	.159	.074	.138	.142	.083	.057	.040	.037	.222
Chinese	Test	–	.007	–	.041	–	.006	–	–	-.006	-.012

like Mandarin, the meaning associated with a syllable depends on its pitch contour. Such language dependent peculiarities make cross-cultural settings highly challenging, especially when noise comes into play because of the ecological conditions of study.

6 CONCLUSIONS

In this paper, we introduced AVEC 2019 – the sixth combined open Audio/Visual Emotion and Health assessment challenge. It comprises three Sub-challenges: i) SoMS, where the level of mood has to be predicted from positive and negative personal stories; ii) DDS, where the level of depression (PHQ-8 score) has to be predicted from structured interviews conducted by a virtual agent wholly driven by AI; and, iii) CES, where the level of affective dimensions of *arousal*, *valence*, and *liking* has to be inferred in a cross-cultural *in-the-wild* paradigm with German and Hungarian cultures as training and testing material, and Chinese culture as solely testing material.

By intention, we opted to use exclusively open-source software and the highest possible transparency and realism for the baselines, by using the same number of trials as given to participants for reporting results on the test partition, and sharing all the developed scripts for both features extraction and machine learning on a public platform. Results indicate that: i) in the SoMS, the level of mood was best predicted when the system was trained on the static scores and evaluated on their dynamic view, i. e., between the label provided after the storytellings, and before the first story,

which can be explained by inertial emotion theories; ii) in the DDS, prediction of the level of depression (PHQ-8) is reported to be more challenging when the virtual agent conducting the interview is wholly driven by AI, compared to a WoZ setup; and iii), in the CES, dimensional emotions are more challenging to sense in a cross-cultural setting for audio descriptors compared to video descriptors, which confirm on one hand the universality of facial expressions for Asian (Chinese) and Western European cultures (German and Hungarian), and show on the other the challenge of using audio descriptors for paralinguistics analysis in languages presenting dissimilarities in their acoustic, in particular when data are collected in an ecological (noisy) environment.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the Horizon 2020 Programme through the Research Innovation Action No. 826506 (sustAGE), and No. 688835 (DE-ENIGMA). Further funding has also been received from the Innovative Medicines Initiative 2 Joint Undertaking under grant agreement No. 115902, which receives support from the European Union’s Horizon 2020 research and innovation program and EFPIA. The work on the DDS was supported in part by the U.S. Army. Any opinion, content or information presented does not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred. The authors further thank the sponsor of the challenge – audEERING GmbH.

REFERENCES

- [1] Tim Althoff, Kevin Clark, and Jure Leskovec. 2016. Large-scale Analysis of Counseling Conversations: An Application of Natural Language Processing to Mental Health. *Transactions of the Association for Computational Linguistics* 4 (2016), 463–476.
- [2] Shahin Amiriparian, Maurice Gerczuk, Sandra Ottl, Nicholas Cummins, Michael Freitag, Sergey Pugachevskiy, Alice Baird, and Björn Schuller. 2017. Snore sound classification using image-based deep spectrum features. In *Proc. of INTERSPEECH 2017, 18th Annual Conference of the International Speech Communication Association*. ISCA, Stockholm, Sweden, 3512–3516.
- [3] American Psychiatric Association. 2013. *Diagnostic and Statistical Manual of Mental Disorders (DSM-5)*. American Psychiatric Publishing, Arlington, VA.
- [4] Tadas Baltrušaitis, Amir Zadeh, Yao Chong Lim, and Louis-Philippe Morency. 2018. OpenFace 2.0: Facial Behavior Analysis Toolkit. In *Proc. 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*. IEEE, Xi'an, P. R. China, 59–66.
- [5] Amit Baumel and Elad Yom-Tov. 2018. Predicting user adherence to behavioral eHealth interventions in the real world: examining which aspects of intervention design matter most. *Translational Behavioral Medicine* 8, 5 (2018), 793–798.
- [6] Yoshua Bengio, Aaron Courville, and Pascal Vincent. 2013. Representation Learning: A Review and New Perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35, 4 (August 2013), 1798–1828.
- [7] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum Learning. In *Proc. International Conference on Machine Learning (ICML)*. ACM, Montreal, QC, Canada, 41–48.
- [8] Jeffrey F. Cohn, Tomas Simon Kruez, Iain Matthews, Ying Yang, Minh Hoai Nguyen, Margara Tejera Padilla, Feng Zhou, and Fernando De la Torre. 2009. Detecting Depression from Facial Actions and Vocal Prosody. In *Proc. 3rd International Conference on Affective Computing and Intelligent Interaction and Workshops*. IEEE, Amsterdam, Netherlands, 7 pages.
- [9] Daniel T. Cordaro, Rui Sun, Dacher Keltner, Shanmukh Kamble, Niranjan Huddar, and Galen McNeil. 2018. Universals and cultural variations in 22 emotional expressions across five cultures. *Emotion* 18 (2018), 75–93.
- [10] Ciprian A. Corneanu, Marc O. Simón, Jeffrey F. Cohn, and Sergio E. Guerrero. 2016. Survey on RGB, 3D, Thermal, and Multimodal Approaches for Facial Expression Recognition: History, Trends, and Affect-Related Applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38, 8 (August 2016), 1548–1568.
- [11] Nicholas Cummins, Stefan Scherer, Jarek Krajewski, Sebastian Schnieder, Julien Epps, and Thomas F. Quatieri. 2015. A review of depression and suicide risk assessment using speech analysis. *Speech Communication* 71 (July 2015), 10–49.
- [12] Kerstin Dautenhahn. 2002. The origins of narrative: In search of the transactional format of narratives in humans and other social animals. *International Journal of Cognition and Technology* 1, 1 (2002), 97–123.
- [13] Jun Deng, Nicholas Cummins, Maximilian Schmitt, Kun Qian, Fabien Ringeval, and Björn Schuller. 2017. Speech-based diagnosis of autism spectrum condition by generative adversarial network representations. In *Proc. 7th International Conference on Digital Health (DH)*. ACM, London, UK, 53–57.
- [14] David DeVault, Ron Artstein, Grace Benn, Teresa Dey, Ed Fast, Alesia Gainer, Kallirroi Georgila, Jonathan Gratch, Arno Hartholt, Margaux Lhommet, Gale Lucas, Stacy Marsella, Fabrizio Morbini, Angela Nazarian, Stefan Scherer, Giota Stratou, Apar Suri, David Traum, Rachel Wood, Yuyu Xu, Alberto Rizzo, and Louis-Philippe Morency. 2014. SimSensei Kiosk: A Virtual Human Interviewer for Healthcare Decision Support. In *Proc. International Conference on Autonomous Agents and Multi-Agent Systems, AAMAS 2014*. ACM, Paris, France, 1061–1068.
- [15] Sidney K. D'Mello and Jacqueline Kory. 2015. A Review and Meta-Analysis of Multimodal Affect Detection Systems. *Comput. Surveys* 47, 3 (February 2015), Article 43, 36 pages.
- [16] Paul Ekman. 1971. Universals and cultural differences in facial expressions of emotion. In *Nebraska Symposium on Motivation*, Vol. 19. University of Nebraska Press, Lincoln, NE, 207–283.
- [17] Hillary Anger Elfenbein and Nalini Ambady. 2002. On the universality and cultural specificity of emotion recognition: A meta-analysis. *Psychological Bulletin* 128, 2 (2002), 203–235.
- [18] Anna Esposito, Antonietta M. Esposito, and Carl Vogel. 2015. Needs and challenges in human computer interaction for processing social emotional information. *Pattern Recognition Letters* 66 (November 2015), 41–51. Issue C.
- [19] Florian Eyben, Klaus R. Scherer, Björn Schuller, Johan Sundberg, Elisabeth André, Carlos Busso, Laurence Devillers, Julien Epps, Petri Laukka, Shrikanth S. Narayanan, and Khiet P. Truong. 2016. The Geneva Minimalistic Acoustic Parameter Set (GeMAPS) for Voice Research and Affective Computing. *IEEE Transactions on Affective Computing* 7, 2 (April 2016), 190–202.
- [20] Florian Eyben, Felix Weninger, Florian Groß, and Björn Schuller. 2013. Recent Developments in openSMILE, the Munich Open-Source Multimedia Feature Extractor. In *Proc. 21st ACM International Conference on Multimedia (ACM MM)*. ACM, Barcelona, Spain, 835–838.
- [21] Silvia Monica Feraru, Dagmar Schuller, and Björn Schuller. 2015. Cross-Language Acoustic Emotion Recognition: An Overview and Some Tendencies. In *Proc. 6th Biannual Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, Xi'an, P. R. China, 125–131.
- [22] Yuan Gong and Christian Poellabauer. 2017. Topic Modeling Based Multi-modal Depression Detection. In *Proc. 7th International Workshop on Audio/Visual Emotion Challenge (AVEC)*. ACM, Mountain View (CA), USA, 69–76.
- [23] Jonathan Gratch, Ron Artstein, Gale Lucas, Giota Stratou, Stefan Scherer, Angela Nazarian, Rachel Wood, Jill Boberg, David DeVault, Stacy Marsella, David Traum, Skip Rizzo, and Louis-Philippe Morency. 2014. The Distress Analysis Interview Corpus of human and computer interviews. In *Proc. 9th International Conference on Language Resources and Evaluation, LREC 2014*. ELRA, Reykjavik, Iceland, 3123–3128.
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, Las Vegas, NV, 770–778.
- [25] Marlies Houben, Wim Van Den Noortgate, and Peter Kuppens. 2015. The relation between short term emotion dynamics and psychological well-being: A meta-analysis. *Psychological Bulletin* 141, 4 (July 2015), 901–930.
- [26] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. 2017. Densely Connected Convolutional Networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Honolulu, HI, 4700–4708.
- [27] Jian Huang, Ya Li, Jianhua Tao, Zheng Lian, Mingyue Niu, and Minghao Yang. 2018. Multimodal Continuous Emotion Recognition with Data Augmentation Using Recurrent Neural Networks. In *Proc. 8th International Workshop on Audio/Visual Emotion Challenge, AVEC'18*. ACM, Seoul, South Korea, 57–64.
- [28] Jyoti Joshi, Roland Goecke, Sharifa Alghowinem, Abhinav Dhall, Michael Wagner, Julien Epps, Gordon Parker, and Michael Breakspear. 2013. Multimodal assistive technologies for depression diagnosis and monitoring. *Journal on Multimodal User Interfaces* 7, 3 (2013), 217–228.
- [29] Heysem Kaya and Alexey A. Karpov. 2018. Efficient and effective strategies for cross-corpus acoustic emotion recognition. *Neurocomputing* 275 (January 2018), 1028–034.
- [30] Dimitrios Kollias, Panagiotis Tzirakis, Mihalis A. Nicolaou, Athanasios Papaioannou, Guoying Zhao, Björn Schuller, Irene Kotsia, and Stefanos Zafeiriou. 2019. Deep affect prediction in-the-wild: Aff-wild database and challenge, deep architectures, and beyond. *International Journal of Computer Vision* 127, 6 (2019), 907–929.
- [31] Jean Kossaifi, Robert Walecki, Yannis Panagakis, Jie Shen, Maximilian Schmitt, Fabien Ringeval, Jing Han, Vedhas Pandit, Björn Schuller, Kam Star, Elnar Hajiye, and Maja Pantic. 2019. SEWA DB: A Rich Database for Audio-Visual Emotion and Sentiment Research in the Wild. <https://arxiv.org/abs/1901.02839>. 17 pages.
- [32] Peter Koval, Peter Kuppens, Nicholas B. Allen, and Lisa Sheeber. 2012. Getting stuck in depression: The roles of rumination and emotional inertia. *Cognition & Emotion* 26, 8 (2012), 1412–1427.
- [33] Peter Koval, Madeline L. Pe, Kristof Meers, and Peter Kuppens. 2013. Affect dynamics in relation to depressive symptoms: Variable, unstable or inert? *Emotion* 13, 6 (2013), 1132.
- [34] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems* 25, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Eds.). Curran Associates, Inc., Lake Tahoe, NV, 1097–1105.
- [35] Peter Kuppens, Nicholas B. Allen, and Lisa B. Sheeber. 2010. Emotional inertia and psychological maladjustment. *Psychological Science* 21, 7 (2010), 984–991.
- [36] Lin Li. 1989. A concordance correlation coefficient to evaluate reproducibility. *Biometrics* 45, 1 (March 1989), 255–268.
- [37] Reza Lotfian and Carlos Busso. 2019. Curriculum Learning for Speech Emotion Recognition from Crowdsourced Labels. *IEEE Transactions on Audio, Speech & Language Processing* 27, 4 (2019), 815–826.
- [38] Humberto R. Maturana and Francisco J. Varela. 1987. *Tree of Knowledge: The Biological Roots of Human Understanding*. New Science Library/Shambhala Publications, Boston, MA.
- [39] Michelle Morales, Stefan Scherer, and Rivka Levitan. 2017. A Cross-modal Review of Indicators for Depression Detection Systems. In *Proc. 4th Workshop on Computational Linguistics and Clinical Psychology – From Linguistic Signal to Clinical Reality*. ACL, Vancouver, BC, 1–12.
- [40] World Health Organization. 2017. *Depression and Other Common Mental Disorders: Global Health Estimates*. Technical Report. World Health Organization. Licence: CC BY-NC-SA 3.0 IGO.
- [41] Vedhas Pandit and Björn Schuller. 2019. On Many-to-Many Mapping Between Concordance Correlation Coefficient and Mean Square Error. <https://arxiv.org/abs/1902.05180>. 23 pages.
- [42] Maja Pantic, Nicu Sebe, Jeffrey F. Cohn, and Thomas Huang. 2005. Affective Multimodal Human-computer Interaction. In *Proc. 13th Annual ACM International Conference on Multimedia*. ACM, Singapore, Singapore, 669–676.
- [43] James W Pennebaker, Matthias R Mehl, and Kate G Niederhoffer. 2003. Psychological Aspects of Natural Language Use: Our Words, Our Selves. *Annual Review of Psychology* 54, 1 (2003), 547–577.
- [44] Nairan Ramirez-Esparza, Cindy K Chung, Ewa Kacwicz, and James W Pennebaker. 2008. The Psychology of Word Use in Depression Forums in English and

- in Spanish: Texting Two Text Analytic Approaches. In *International Conference on Weblogs and Social Media*. AAAI, Seattle, WA, 102–108.
- [45] Eva-Maria Rathner, Julia Djamali, Yannik Terhorst, Björn Schuller, Nicholas Cummins, Gudrun Salamon, Christina Hunger-Schoppe, and Harald Baumeister. 2018. How Did You like 2017? Detection of Language Markers of Depression and Narcissism in Personal Narratives. In *Proc. of INTERSPEECH 2018, 19th Annual Conference of the International Speech Communication Association*. ISCA, Hyderabad, India, 3388–3392.
- [46] Eva-Maria Rathner, Yannik Terhorst, Nicholas Cummins, Björn Schuller, and Harald Baumeister. 2018. State of Mind: Classification through Self-reported Affect and Word Use in Speech. In *Proc. of INTERSPEECH 2018, 19th Annual Conference of the International Speech Communication Association*. ISCA, Hyderabad, India, 267–271.
- [47] Fabien Ringeval, Björn Schuller, Michel Valstar, Roddy Cowie, Heysem Kaya, Maximilian Schmitt, Shahin Amiriparian, Nicholas Cummins, Dennis Lalanne, Adrien Michaud, Elvan Ciftci, Hüseyin Gülec, Albert Ali Salah, and Maja Pantic. 2018. AVEC 2018 Workshop and Challenge: Bipolar Disorder and Cross-Cultural Affect Recognition. In *Proc. 8th International Workshop on Audio/Visual Emotion Challenge, AVEC'18*. ACM, Seoul, South Korea, 3–13.
- [48] Fabien Ringeval, Björn Schuller, Michel Valstar, Roddy Cowie, and Maja Pantic. 2015. AVEC 2015 – The 5th International Audio/Visual Emotion Challenge and Workshop. In *Proc. 23rd ACM International Conference on Multimedia, MM 2015*. ACM, Brisbane, Australia, 1335–1336.
- [49] Fabien Ringeval, Björn Schuller, Michel Valstar, Roddy Cowie, and Maja Pantic. 2017. Summary for AVEC 2017 – Real-life Depression, and Affect Recognition Challenge and Workshop. In *Proc. 25th ACM International Conference on Multimedia (ACM MM)*. ACM, Mountain View, CA, USA, 1963–1964.
- [50] Fabien Ringeval, Björn Schuller, Michel Valstar, Roddy Cowie, and Maja Pantic. 2018. Summary for AVEC 2018: Bipolar Disorder and Cross-Cultural Affect Recognition. In *Proc. 26th ACM International Conference on Multimedia, MM 2018*. ACM, Seoul, South Korea, 2111–2112.
- [51] Fabien Ringeval, Björn Schuller, Michel Valstar, Jonathan Gratch, Roddy Cowie, Stefan Scherer, Sharon Mozgai, Nicholas Cummins, and Maja Pantic. 2017. AVEC 2017 – Real-life Depression, and Affect Recognition Workshop and Challenge. In *Proc. 7th International Workshop on Audio/Visual Emotion Challenge (AVEC)*. ACM, Mountain View, CA, USA, 3–9.
- [52] James A Russell. 2003. Core affect and the psychological construction of emotion. *Psychological review* 110, 1 (2003), 145.
- [53] James A Russell and Ulrich F Lanius. 1984. Adaptation level and the affective appraisal of environments. *Journal of Environmental Psychology* 4, 2 (1984), 119–135.
- [54] Hesam Sagha, Jun Deng, Maryna Gavryukova, Jing Han, and Björn Schuller. 2016. Cross lingual speech emotion recognition using canonical correlation analysis on principal component subspace. In *Proc. 41st IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, Shanghai, P. R. China, 5800–5804.
- [55] Robert M Sapolsky. 2004. Social status and health in humans and other animals. *Annu. Rev. Anthropol.* 33 (2004), 393–418.
- [56] Klaus R. Scherer, Rainer Banse, and Harald G. Wallbott. 2001. Emotion inferences from vocal expression correlate across languages and cultures. *Journal of Cross-Cultural Psychology* 32, 1 (January 2001), 76–92.
- [57] Stefan Scherer, Giota Stratou, Jonathan Gratch, Jill Boberg, Marwa Mahmoud, Albert (Skip) Rizzo, and Louis-Philippe Morency. 2013. Automatic Behavior Descriptors for Psychological Disorder Analysis. In *Proc. 10th IEEE International Conference and Workshops on Automatic Face & Gesture Recognition (FG)*. IEEE, Shanghai, P. R. China, 8 pages.
- [58] Stefan Scherer, Giota Stratou, Gale Lucas, Marwa Mahmoud, Jill Boberg, Jonathan Gratch, Albert (Skip) Rizzo, and Louis-Philippe Morency. 2014. Automatic audio-visual behavior descriptors for psychological disorder analysis. *Image and Vision Computing* 32, 10 (October 2014), 648–658.
- [59] Maximilian Schmitt, Fabien Ringeval, and Björn Schuller. 2016. At the border of acoustics and linguistics: Bag-of-Audio-Words for the recognition of emotions in speech. In *Proc. of INTERSPEECH 2016, 17th Annual Conference of the International Speech Communication Association*. ISCA, San Francisco, CA, USA, 495–499.
- [60] Maximilian Schmitt and Björn Schuller. 2017. openXBOW – Introducing the Passau Open-Source Crossmodal Bag-of-Words Toolkit. *Journal of Machine Learning Research* 18, 96 (2017), 1–5.
- [61] Björn Schuller, Stefan Steidl, Anton Batliner, Peter B. Marschik, Harald Baumeister, Fengquan Dong, Simone Hantke, Florian B. Pokorny, Eva-Maria Rathner, Katrin D. Bartl-Pokorny, Christa Einspieler, Dajie Zhang, Alice Baird, Shahin Amiriparian, Kun Qian, Zhao Ren, Maximilian Schmitt, Panagiotis Tzirakis, and Stefanos Zafeiriou. 2018. The INTERSPEECH 2018 Computational Paralinguistics Challenge: Atypical & Self-Assessed Affect, Crying & Heart Beats. In *Proc. of INTERSPEECH 2018, 19th Annual Conference of the International Speech Communication Association*. ISCA, Hyderabad, India, 122–126.
- [62] Björn Schuller, Michel Valstar, Florian Eyben, Roddy Cowie, and Maja Pantic. 2012. AVEC 2012 – The continuous Audio/Visual Emotion Challenge. In *Proc. 14th ACM International Conference on Multimodal Interaction (ICMI)*. ACM, Santa Monica, CA, USA, 449–456.
- [63] Björn Schuller, Michel Valstar, Florian Eyben, Gary McKeown, Roddy Cowie, and Maja Pantic. 2011. AVEC 2011 – The First International Audio/Visual Emotion Challenge. In *Proc. 4th Biannual International Conference on Affective Computing and Intelligent Interaction (ACII)*, Vol. II. Springer, Memphis, TN, USA, 415–424.
- [64] Norbert Schwarz and Gerard L. Clore. 1983. Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states. *Journal of Personality and Social Psychology* 45, 3 (September 1983), 512–523.
- [65] Andreas Schwerdtfeger. 2004. Predicting autonomic reactivity to public speaking: don't get fixed on self-report data! *International Journal of Psychophysiology* 52, 3 (2004), 217–224.
- [66] Andreas R Schwerdtfeger and Eva-Maria Rathner. 2016. The ecological validity of the autonomic-subjective response dissociation in repressive coping. *Anxiety, Stress, & Coping* 29, 3 (2016), 241–258.
- [67] Caifeng Shan, Shaogang Gong, and Peter W Mcowan. 2009. Facial expression recognition based on Local Binary Patterns: A comprehensive study. *Image and Vision Computing* 27, 6 (2009), 803–816.
- [68] Stewart Shapiro and Deborah J. MacInnis. 2002. Understanding program-induced mood effects: Decoupling arousal from valence. *Journal of Advertising* 31, 4 (May 2002), 15–26.
- [69] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. <https://arxiv.org/abs/1409.1556>. 14 pages.
- [70] Lukas Stappen, Nicholas Cummins, Eva Messner, Harald Baumeister, Judith Dineley, and Björn Schuller. 2019. Context Modelling Using Hierarchical Attention Networks for Sentiment and Self-assessed Emotion Detection in Spoken Narratives. In *Proc. 44th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, Brighton, United Kingdom, 6680–6684.
- [71] George Trigeorgis, Fabien Ringeval, Raymond Brueckner, Erik Marchi, Mihalis A. Nicolaou, Björn Schuller, and Stefanos Zafeiriou. 2016. Adieu features? End-to-end speech emotion recognition using a deep Convolutional Recurrent Network. In *Proc. 41st IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, Shanghai, P. R. China, 5200–5204.
- [72] Michel Valstar, Jonathan Gratch, Björn Schuller, Fabien Ringeval, Roddy Cowie, and Maja Pantic. 2016. Summary for AVEC 2016: Depression, mood, and emotion recognition workshop and challenge. In *Proc. 24th ACM International Conference on Multimedia (ACM MM)*. ACM, Amsterdam, The Netherlands, 1483–1484.
- [73] Michel Valstar, Jonathan Gratch, Björn Schuller, Fabien Ringeval, Denis Lalanne, Mercedes Torres Torres, Stefan Scherer, Giota Stratou, Roddy Cowie, and Maja Pantic. 2016. AVEC 2016 – Depression, mood, and emotion recognition workshop and challenge. In *Proc. 6th International Workshop on Audio/Visual Emotion Challenge (AVEC)*. ACM, Amsterdam, The Netherlands, 3–10.
- [74] Michel Valstar, Björn Schuller, Jarek Krajewski, Roddy Cowie, and Maja Pantic. 2013. Workshop summary for the 3rd international Audio/Visual Emotion Challenge and workshop. In *Proc. 21st ACM International Conference on Multimedia (ACM MM)*. ACM, Barcelona, Spain, 1085–1086.
- [75] Michel Valstar, Björn Schuller, Jarek Krajewski, Roddy Cowie, and Maja Pantic. 2014. AVEC 2014: The 4th international Audio/Visual Emotion Challenge and workshop. In *Proc. 22nd ACM International Conference on Multimedia (ACM MM)*. ACM, Orlando, FL, USA, 1243–1244.
- [76] Kalani Wataraka Gamage, Ting Dang, Vidhyasaharan Sethu, Julien Epps, and Eliathamby Ambikairajah. 2018. Speech-based Continuous Emotion Prediction by Learning Perception Responses Related to Salient Events: A Study Based on Vocal Affect Bursts and Cross-Cultural Affect in AVEC 2018. In *Proc. 8th International Workshop on Audio/Visual Emotion Challenge, AVEC'18*. ACM, Seoul, South Korea, 47–55.
- [77] Felix Weninger, Fabien Ringeval, Erik Marchi, and Björn Schuller. 2016. Discriminatively trained recurrent neural networks for continuous dimensional emotion recognition from audio. In *Proc. 25th International Joint Conference on Artificial Intelligence (IJCAI)*. IJCAI/AAAI, New York City, NY, USA, 2196–2202.
- [78] James R. Williamson, Thomas F. Quatieri, Brian S. Helfer, Rachele Horwitz, Bea Yu, and Daryush D. Mehta. 2013. Vocal Biomarkers of Depression Based on Motor Incoordination. In *Proc. 3rd ACM International Workshop on Audio/Visual Emotion Challenge (AVEC'13)*. ACM, Barcelona, Spain, 41–48.
- [79] Georgios N. Yannakakis, Roddy Cowie, and Carlos Busso. 2017. The ordinal nature of emotions. In *Proc. 7th Biannual Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, San Antonio, TX, 248–255.
- [80] Biquiao Zhang, Emily Mower Provost, and Georg Essl. 2017. Cross-corpus acoustic emotion recognition with multi-task learning: Seeking common ground while preserving differences. *IEEE Transactions on Affective Computing* 10, 1 (March 2017), 85–99.
- [81] Zixing Zhang, Nicholas Cummins, and Björn Schuller. 2017. Advanced data exploitation in speech analysis – An overview. *IEEE Signal Processing Magazine* 34, 4 (July 2017), 107–129.
- [82] Jiming Zhao, Ruichen Li, Shizhe Chen, and Qin Jin. 2018. Multi-modal Multi-cultural Dimensional Continues Emotion Recognition in Dyadic Interactions. In *Proc. 8th International Workshop on Audio/Visual Emotion Challenge, AVEC'18*. ACM, Seoul, South Korea, 65–72.