



MuSe 2020 Challenge and Workshop: Multimodal Sentiment Analysis, Emotion-target Engagement and Trustworthiness Detection in Real-life Media

Emotional Car Reviews in-the-wild

Lukas Stappen
University of Augsburg
Augsburg, Germany

Alice Baird
University of Augsburg
Augsburg, Germany

Georgios Rizos
Imperial College London
London, UK

Panagiotis Tzirakis
Imperial College London
London, UK

Xinchen Du
Technical University of Munich
Munich, Germany

Felix Hafner
University of Augsburg
Augsburg, Germany

Lea Schumann
University of Augsburg
Augsburg, Germany

Adria Mallol-Ragolta
University of Augsburg
Augsburg, Germany

Björn W. Schuller
Imperial College London
London, United Kingdom

Iulia Lefter
Delft University of Technology
Delft, Netherlands

Erik Cambria
Nanyang Technological University
Singapore

Ioannis Kompatsiaris
CERTH - ITI
Thermi-Thessaloniki, Greece

ABSTRACT

Multimodal Sentiment Analysis in Real-life Media (MuSe) 2020 is a Challenge-based Workshop focusing on the tasks of sentiment recognition, as well as emotion-target engagement and trustworthiness detection by means of more comprehensively integrating the audio-visual and language modalities. The purpose of MuSe 2020 is to bring together communities from different disciplines; mainly, the audio-visual emotion recognition community (signal-based), and the sentiment analysis community (symbol-based). We present three distinct sub-challenges: **MuSe-Wild**, which focuses on continuous emotion (arousal and valence) prediction; **MuSe-Topic**, in which participants recognise 10 domain-specific topics as the target of 3-class (low, medium, high) emotions; and **MuSe-Trust**, in which the novel aspect of trustworthiness is to be predicted. In this paper, we provide detailed information on **MuSe-CaR**, the first of its kind in-the-wild database, which is utilised for the challenge, as well as the state-of-the-art features and modelling approaches applied. For each sub-challenge, a competitive baseline for participants is set; namely, on test we report for **MuSe-Wild** a combined (valence and arousal) CCC of .2568, for **MuSe-Topic** a score (computed as $0.34 \cdot \text{UAR} + 0.66 \cdot \text{F1}$) of 76.78 % on the 10-class topic and 40.64 % on the 3-class emotion prediction, and for **MuSe-Trust** a CCC of .4359.

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, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MuSe'20, October 16, 2020, Seattle, WA, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8157-4/20/10...\$15.00
<https://doi.org/10.1145/3423327.3423673>

CCS CONCEPTS

• **Information systems** → **Multimedia and multimodal retrieval**; • **Computing methodologies** → **Artificial intelligence**.
KEYWORDS

Multimodal Sentiment Analysis; Affective Computing; User-Generated Data; Multimodal Fusion

ACM Reference Format:

Lukas Stappen, Alice Baird, Georgios Rizos, Panagiotis Tzirakis, Xinchen Du, Felix Hafner, Lea Schumann, Adria Mallol-Ragolta, Björn W. Schuller, Iulia Lefter, Erik Cambria, and Ioannis Kompatsiaris. 2020. MuSe 2020 Challenge and Workshop: Multimodal Sentiment Analysis, Emotion-target Engagement and Trustworthiness Detection in Real-life Media: Emotional Car Reviews in-the-wild. In *1st International Multimodal Sentiment Analysis in Real-life Media Challenge and Workshop (MuSe'20), October 16, 2020, Seattle, WA, USA*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3423327.3423673>

1 INTRODUCTION

Multimodal Sentiment Analysis in Real-life Media (MuSe) 2020 is a novel Challenge-based Workshop in which *sentiment recognition*, as well as *emotion-target engagement* and *trustworthiness detection* are the main focus. MuSe aims to provide a testing bed for more extensively exploring the fusion of the audio-visual and language modalities. The core purpose of MuSe is to bring together communities from differing computational disciplines; mainly, the sentiment analysis community (symbol-based), and the audio-visual emotion recognition community (signal-based).

The first group – rooted in the field of Sentiment (and Opinion) Mining and specialising in Natural Language Processing (NLP) methods for symbolic information analysis – leverages the text modality, and focuses on the prediction only of discrete sentiment label categories [50]. In numerous competitions from recent years

researchers from the second group – mostly rooted in the field of Affective (and Behavioural) Computing and specialised in intelligent signal processing – focused on one, or both of the audio and vision modalities, in order to predict the continuous-valued valence and arousal dimensions of emotion (circumplex model of affect), while often disregarding the potential contribution of textual information [15, 30, 36, 47]. However, approaches by both communities now show signs of convergence, highly influenced by related, explicitly multimodal learning techniques [3, 13, 28]. Of note, the 2020 INTERSPEECH Computational Paralinguistics (ComParE) Challenge included for the first time baselines utilising both audio signal and text transcripts [35].

With this in mind, MuSe 2020 aims to attract both communities equally and encourages a fusion of modalities to demonstrate the advantages within the field of emotion specifically. Ideally, participation should strive towards the development of unified approaches applicable to each task. Tasks have arisen from different academic traditions: on the one hand, complex, dimensional emotion annotations relating to the expression of behaviour, and on the other hand, linking sentiment and emotions to topics (context), entities or aspects, as is common in sentiment analysis [38].

A second contribution of MuSe 2020 is the facilitation of a broad comparison of the merits for the three core modalities (language, audio, and visual cues), as well as various approaches of multimodal fusion under well-defined and strictly comparable conditions. In this way, establishing the extent to which the fusion of approaches is possible and beneficial, as well as advancing sentiment and emotion recognition systems to be able to deal with fully naturalistic (in-the-wild) behaviour from large volumes of in-the-wild (user-generated) data. User-generated data types refers to data sourced from the target user themselves and are the new generation of data utilised for real world multimedia affect and sentiment analysis [48] and other research fields [9].

For all of the three sub-challenges, one dataset is chosen to make the comparison between each sub-challenge more easily facilitated. For this year’s MuSe 2020, we introduce the Multimodal Sentiment Analysis in Car Reviews dataset *MuSe-CaR* which covers the range of aforementioned topics discussed. MuSe-CaR is a large, multimodal dataset which has been gathered in-the-wild with the intention of further understanding real world Multimodal Sentiment Analysis, in particular the emotional engagement that takes place during product reviews (i. e., automobile reviews) where a sentiment is linked to a topic or entity.

2 CHALLENGE OUTLINE AND PROTOCOL

The major novelties discussed herein will be introduced in MuSe 2020 through three core sub-challenges, (i) Multimodal Sentiment in-the-Wild Sub-challenge (MuSe-Wild), (ii) Multimodal Emotion-Target Engagement Sub-challenge (MuSe-Topic) (iii) Multimodal Trustworthiness Sub-challenge (MuSe-Trust). In the following, we will describe and highlight the aforementioned novelties of each each sub-challenge, as well as include guidelines for participation.

Individuals wishing to participate in the MuSe 2020 challenge must hold an academic affiliation. Further to this, they should download and fill out the End User License Agreement (EULA) and submit

via the homepage¹. All entries to the challenge should be accompanied with a document which describes in detail methods and results and includes a citation of this paper. To appear on the temporary, public leader board on the MuSe homepage, participants must provide predictions, a Github repository where their source code is uploaded, and a link to an arXiv preliminary technical report. The organisers do not participate in the Challenge themselves, but re-evaluate the findings of the best performing system of each Sub-challenge. There will be a double blind peer-reviewed process by the technical program committee, and only papers which meet the standards set by peer-review will be eligible for the main competition. Papers accepted for the workshop will be allocated 6-8 pages (plus references) in the proceedings of ACM MM 2020.

2.1 MuSe-Wild Sub-Challenge

In the *MuSe-Wild Sub-Challenge*, participants are predicting the level of affective dimensions (arousal, and valence) in a time-continuous manner from audio-visual recordings. Valence thereby is strongly linked to the emotional component of the umbrella term of sentiment analysis and is often used interchangeably [22, 27, 41]. Timestamps to enable modality alignment and fusion on word-, sentence-, and utterance-level as well as several acoustic, visual and textual-based features are pre-computed and provided with the challenge package. The evaluation metric for this sub-challenge is *concordance correlation coefficient (CCC)*, which is often used in similar challenges [30, 47]. CCC is a measure of reproducibility and performance, which condenses information on both precision and accuracy, is robust to changes in scale and location [18], and its theoretical properties to other regression measures, e. g., (root) mean squared error, are well understood [23]. For the baseline for the MuSe-Wild sub-challenge the mean of arousal and valence is taken.

2.2 MuSe-Topic Sub-challenge

In the *MuSe-Topic Sub-challenge*, participants are predicting 10-classes of domain-specific (automotive, as given by the chosen database) topics² as the target of emotions. In addition, three classes (low, medium, and high) of valence and arousal should be predicted i. e., for each topic segment, one valence and one arousal value. These classes are the mean value of the temporally aggregated continuous labels of MuSe-Wild, which were divided into three equally sized classes (33 %) for each label. For this sub-challenge, first, the weighted score combining (0.34·) Unweighted Average Recall (UAR) and (0.66·) F1 (micro) measures independently for each predictions (Valence, Arousal and Topic) are calculated. We include both these factors to keep our evaluation consistent with previous challenges, as the former was partially used to evaluate a classification task in [15], and the latter in [35]. Second, the mean of the weighted scores for Valence and Arousal (combined) is calculated. Third, to combine the mean with the topic score the mean rank over all participants ((rank of combined emotions result + rank of topic result)/2) is calculated for the final performance assessment. In case two participants should have the same mean rank, the one with the

¹www.muse-challenge.org

²Classes for the MuSe-Topic sub-challenge include: General Information, Costs, Performance, Quality & Aesthetic, Safety, Comfort, Exterior Features, Interior Features, Handling/Driving Experience, User Experience.

Table 1: Partitioning of the MuSe-CaR dataset, applied for each of the three sub-challenges. Reported are the number of unique videos, and the duration for each sub-challenge hh :mm ss. The unprocessed duration of the MuSe-CaR dataset is 36 :52 :08.

Partition	No.	MuSe-Wild	MuSe-Topic	MuSe-Trust
Train	166	22:16:43	22:35:55	22:45:52
Devel.	62	06:48:58	06:49:46	06:52:22
Test	64	06:02:20	06:14:08	06:12:53
Σ	291	35:08:01	35:39:49	35:51:07

highest topic rank will be the final winner. We believe that this composite measure is most discriminative to meaningfully showcase performance improvements in emotion and topic prediction, as it places importance on precision and recall, in both a dataset-wide and class-specific manner.

2.3 MuSe-Trust Sub-challenge

In the *MuSe-Trust Sub-challenge*, participants are predicting a continuous trustworthiness signal from user-generated audio-visual content in a sequential manner and are provided with aligned valence and arousal annotations, which participants are encouraged to explore, in a means of understanding the relationship between emotional labels in depth and at large scale. The evaluation metric for this sub-challenge is concordance correlation coefficient (CCC).

3 CHALLENGE DATASET

For all of the three Sub-Challenges of MuSe 2020, the MuSe-CaR data set is utilised. MuSe-CaR is a large, extensively annotated multimodal ((spoken) language, audio, video) dataset which has been gathered in-the-wild with the intention of developing appropriate methods and further understanding Multimodal Sentiment Analysis in-the-wild. MuSe-CaR has been designed with an abundance of computational tasks in mind, including emotion and entity recognition, and dominantly with the intention of improving machine understanding of how sentiment (i. e., emotion) is linked to an entity and aspects of such reviews.

The estimated age range of the professional, semi-professional ('influencers'), and casual reviewers is from the mid-20s until the late-50s. Most are native English speakers from the UK or the US, while a small minority are non-native, yet fluent English speakers. MuSe-CaR includes high voice and video quality, as everyday recording devices have improved in recent years. This enables robust learning of a high degree of novel, in-the-wild characteristics.

For the MuSe 2020 Challenge, we selected a high-quality sub-set of the MuSe-CaR dataset consisting of 36 h : 52 m : 08 s of video data from 291 videos and 70 host speakers (plus an additional of roughly 20 narrators) sourced from YouTube.

When creating the data set, it was of particular importance to find a balance between the stable and uncontrollable, 'in-the-wild' properties such as different recording devices, camera perspectives, ambient noises (car noises, music), or changing backgrounds to allow for meaningful learning with current deep learning methods. Such 'in-the-wild' characteristics of MuSe-CaR include; i) *video*: shot size, face-angle, camera motion, reviewer visibility, reviewer face occlusion (glasses), and highly varying backgrounds; ii) *audio*: ambient noises, narrator and host diarisation, diverse microphone

types, and speaker locations; iii) *text*: colloquialisms, and domain-specific terms.

The topic of videos within MuSe-CaR is limited to vehicle reviews, with the number of vehicle manufacturers being restricted to premium brands (BMW, Audi, Mercedes-Benz) that equip their vehicles with the latest technology, thus, ensuring that discussed entities and aspects (e. g., semi-autonomous vehicle functions) occur across a board range of videos (and different manufacturers). Most of the reviewers are semi- or professional reviewers (e. g., YouTube channel 'influencers'). All YouTube channels used within MuSe-CaR have given full consent for their data to be used with the context of academic research³.

To avoid extremely objective reviews, during the selection process videos were rated on a scale between 0 (emotionless) and 5 (very emotional). We filtered out all videos with a score less than 3 before annotation began. Within MuSe-CaR, there are 15 annotation tiers (3 continuous dimensional, 3 partially continuous binary label, 5 categorical, and 4 automatically annotated tiers). For MuSe 2020, we utilise 3 continuous ratings, and the topic categorical dimensions. Each recording has been annotated in three continuous dimensions; emotional *valence* (hence reflecting sentiment) and *arousal* according to Russell's theory [32], and additionally the novel aspect of *trustworthiness*, each by at least 5 independent annotators. In the case of the Trustworthiness dimension, there has been minimal research into the link between this and other emotions [1], and to the best of the authors' knowledge, it has not been utilised nor predicted using machine learning.

A *gold-standard* was computed on the individual annotators using an Evaluator Weighted Estimator (EWE) approach, in which inter-rater agreement is considered. EWE is described, e. g., further in [34] and has been applied to similar continuous emotion-based tasks [30], and corpora [31]. In addition to the dimensional annotations we included the categorical labelling of emotional engagement with topics, such as comfort, safety, interior, and performance.

For the MuSe 2020 Challenge, data has been partitioned in a Train, Development, and Test convention, where aspects including emotional ratings, speaker independence, and duration have been considered (cf. Table 1 for an overview). The total duration of data for each sub-challenge varies, as further pre-processing to include the most informative data only was applied. For MuSe-Wild and MuSe-Trust, all parts with an active voice or a visible face are included. We excluded non-product related video segments (e. g., advertisements) for MuSe-Wild and MuSe-Topic to minimise the distortion these could cause on the task objectives. More specifically, for MuSe-Topic, we only included sections which have an active voice based on the sentence transcriptions. To not fragment it to purely sentence segments, we fused adjacent segments if the segments cover the same topic and are less than two seconds apart.

³Following the YouTube guidelines, uploading a video to YouTube automatically issues that video under the YouTube own license. To the best of our understanding, under this licence, the use of the data in the EU is only possible by YouTube directly or with the consent of the creator. In similar works, the database producers refer to the fair use principle for academic use. Some videos are also available under the Creative Commons Licence (CC-BY, full use, if the creator credits are mentioned which are provided in the data packages.)

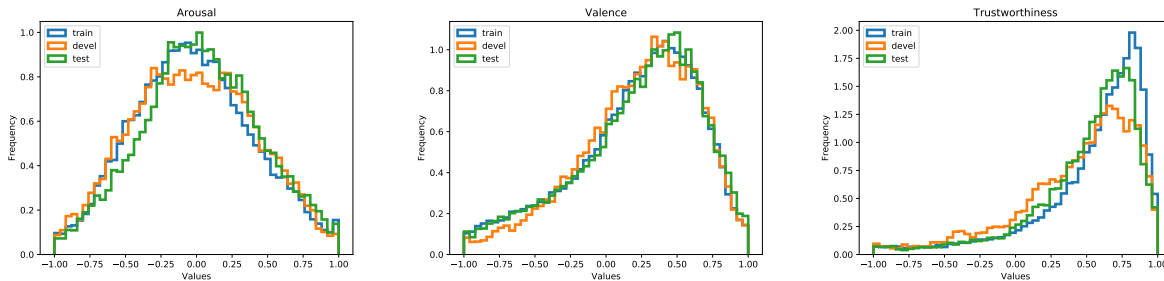


Figure 1: Frequency distribution in the partitions Train, (Devel)opment, and Test for the continuous prediction sub-challenges MuSe-Wild, Arousal and Valence, and MuSe-Trust, Trustworthiness.

Regarding MuSe-Trust, non-product related information – for instance, advertisement –, might have a notable impact on the trustworthiness perception of the video. Therefore, segments containing advertisements for products and YouTube channels are included.

4 BASELINE FEATURES

For each Sub-challenge, we provide a selection of features to participants which have been extracted from language, audio (including speech-to-text), and video signals. Extracting rich features from a huge amount of video data takes days, sometimes weeks, to complete, which would cost participants valuable time. For this reason, we provide 14 model-ready audio, visual, and linguistic feature sets⁴, an amount which far exceeds the number of feature sets provided by other comparable audio-visual challenges [16, 30, 47, 50]. In the proceeding section, the feature sets based on each modality (acoustic, vision, and language) are described. For all feature sets a hop size of 0.25 s was applied (unless otherwise stated) to be inline with the annotation sampling rate.

4.1 Acoustic

For extracting acoustic features, we utilise well-known feature extraction tools, namely `OPENSIMILE` and `DEEPSPECTRUM`, which have both shown success in a variety of audio processing tasks, including prominent work in speech emotion recognition (SER) [8, 33]. Audio is extracted directly from the YouTube videos, normalised to -3dB and converted from stereo to mono, 16 kHz, 16 bit. For all acoustic features, we apply a window size of 5 seconds.

4.1.1 `OPENSIMILE`. The freely available `OPENSIMILE` toolkit [12] is utilised to extract the well-known extended Geneva Minimalistic Acoustic Parameter Set (`eGEMAPS`) [11]. `eGEMAPS` is a hand-crafted speech-based feature set, containing 88 features designed specially for Speech Emotion Recognition (SER) tasks [40]. In addition, 130 dimensional low-level descriptors are provided which have been computed with `OPENSIMILE`, and include the features 1st and 2nd-order derivatives (deltas and double-deltas). LLD extraction remained at the the default `OPENSIMILE` configuration and therefore a window size of 10 ms is applied for this feature set only.

4.1.2 `DEEPSPECTRUM`. We also include `DEEPSPECTRUM` features as a state-of-the-art deep learning based approach [2]. `DEEPSPECTRUM` features extract spectral images from speech instances and

are then fed into pre-trained image recognition Convolutional Neural Networks (CNNs), and the resulting activations are extracted as feature vectors. For MuSe 2020, we extract features utilising the VGG-19 *extraction network* [37], with all other parameters remaining as default. This results in a feature set of 4 096-dimensions.

4.2 Vision

Most visual feature extractors are either designed to localise and extract specific image characteristics and sections (e. g., face), or to learn general discriminatory features for classifying (multi-class, multi-label) a large number of images into many classes (ImageNet). We provide participants with raw data (extracted faces), features focusing on human behavior (face, poses) as well as feature sets which capture the environment as a whole (`XCEPTION`) or the interaction object car (`GoCAR`).

4.2.1 `MTCNN`. To extract and localise the faces in the videos, `MTCNN` [51] was used. Internally, it has a cascaded structure of three stages to predict face and landmark position operating in real time. The model is trained on the data sets `WIDER FACE` [49] and `CelebA` [20]. It also provides a confidence measure which allows the false positive to false negative rate to be tuned. Because the frameworks that extract more detailed face features do not provide features for false positives, we chose not to tune the confidence threshold. For the quantitative performance analysis, we labelled a small selection of videos from each channel by hand, and calculated the intersection over union. Depending on the size of the overlap and intersection, we classified the detected bounding boxes into true and false positives. The detector achieved an accuracy of 90%, and an F1 score of 86% on the selection of MuSe-CaR. In addition, we visually inspected the bounding boxes to control the qualitative performance. Both performances underline the very good quality of `MTCNN` face extractions. These extractions were used as inputs for `VGGFACE` and `OPENFACE`.

4.2.2 `VGGFACE`. `VGGFACE` [24] is used to extract facial features from the cropped faces that were detected by `MTCNN`. Originally intended for face recognition tasks, it outputs a feature vector of size 512 when the top layer is removed. Its main advantage is the comparable performance to other face recognition models while using less data for training. The data set used to train the deep CNN, called `VGG16` [37], is `VGGFACE`, collected by the visual geometry group of Oxford. It contains more than 2 500 identities and 2.6 million faces. While consisting of fewer identities and pose/age variations in comparison to its successor [6], the number of images

⁴Note: The participants are also free to use external data, any commercial or academic feature extractors, pre-trained networks, and libraries. However, this should be reproducible and clearly discussed in the accompanying paper.

is similar in scale. Compared to OpenFace, these features can be used to extract more raw facial features, e. g., to learn predictive facial movements from scratch.

4.2.3 OpenFace. Facial features were also extracted from the cropped faces detected with MTCNN using OPENFACE [4]. This toolkit provides a wide range of facial features. We extracted facial landmarks in both 2D (136 features) and 3D (204 features), 6 head pose features, 288 gaze positions, and the intensity and presence of 17 Facial Action Units (FAUs) each for the left side and centre.

4.2.4 XCEPTION. We use XCEPTION [14] to provide features that capture the environment. XCEPTION is a very deep, state of the art network using residual blocks which enable easier optimisation for large networks. This architecture won the 1st place on the ILSVRC 2015 classification task and other challenges. It is commonly used as feature extractor for general vision features. To obtain the deep representations, we extract the output of the last fully connected layer from the pre-trained XCEPTION network. As a result, a 2048-dimensional deep feature vector is provided for each frame.

4.2.5 GoCAR. GoCAR [39] is an domain-specific visual feature extractor enabling the localisation of 28 car parts, such as, door, steering wheel, headlights, and infotainment with which the reviewer interacts inside and outside the vehicle. It is based on a modified YoloV3 framework [29] with a Darknet-53 as backbone and is trained with a multi-label, multi-class real-world data set containing 15 003 vehicle images of 18 different BMW models with up to 100 different feature variants, each. The coverage of a high number of feature variants is necessary to learn robust features, since cars have one of the highest possible number of product variants, e. g., the number of Mercedes E-Class equipment variations exceeds the order of 10^{24} [26]. The extractor achieves a mean average precision of 67.57 % ranging from 94 % for very distinctive parts such as grills to 14 % for less distinctive ones (e. g., roof window) on 1 000 extracted and manually labelled MuSe-CaR video frames. The provided GoCAR features are converted into an array of fixed size. For this purpose, we use the 10 objects with the highest confidence, and for each object we store the class (one-hot encoded), the confidence and the localisation coordinates (x, y, width, and height). In total, this results in a feature vector of $10 * (27 + 7)$ -dimensions.

4.2.6 OPENPOSE. We extracted 18 2D pose keypoints⁵ using the method proposed in [7], which yielded the best results in the COCO 2016 keypoints challenge [19]. We assume that at maximum, only one person is present in each frame. We use the pre-trained model provided by the authors in [7], trained on the COCO 2016 dataset [19]. The model consists of two branches of stacked CNNs, where one predicts 2D confidence maps for the keypoints of interest, and the other predicts Part Affinity Fields that contain information on the association of keypoints of the same individual amongst themselves. At each level, the outputs of each branch are concatenated and given as input to the higher level layer pair. In the end, we provide the 2D coordinates, as well as the corresponding confidence value of a keypoint being present, for each of the 18 keypoints.

⁵The keypoints are: Nose, Neck, Right/Left Shoulder, Right/Left Elbow, Right/Left Wrist, Right/Left Hip, Right/Left Knee, Right/Left Ankle, Right/Left Eye, and Right/Left Ear.

4.3 Language

FASTTEXT [5] is a library for efficient learning of word embeddings. It is based on the skipgram model where a vector representation is associated to each character n-gram. The model is trained on the English Common Crawl corpus (600B tokens). In comparison to other traditional word embeddings, such as, word2vec [21], or GloVe [25], these sub-words chunks make it possible to calculate word representations of words which were not part of the original training corpus (out-of-vocabulary). This appears advantageous since we work with a domain-specific corpus including technical terms and model names. This a valuable function, and enables us to transform 96 % of words to word embedding vectors.

4.4 Alignment

The wide diversity of feature types from three modalities and the correspondingly different sampling rates lead to different lengths of the extracted features along the time axis. All continuous visual feature extractors (e. g., XCEPTION, GoCAR) sample 4 frames per second, which corresponds approximately to the 250 ms labeling and the 250 ms audio sampling of DEEPSPECTRUM and EGEMAPS (except low-level descriptors which are sampled every 10 ms). Furthermore, Human-focused features (e. g., VGGFACE, FACIAL ACTION UNITS) are only extracted from the frames when a reviewer is visible. Recent work [42] has shown that even when advanced alignment mechanisms are in-cooperated in a multimodal neural network, such as attention heads models, the nets are more effective when the features are first aligned during pre-processing. Therefore, we provide for each sub-challenge non-aligned, label-aligned, and, additionally for the more text-related task MuSe-Topic, FASTTEXT-aligned features. If desired, the non-aligned features can be aligned by the participants using the corresponding timestamps (or start and end time of a segment for MuSe-Topic). The label-aligned features have exactly the same length (and timestamps) as the provided label files. We applied zero-padding to the frames, where the feature type is not present or which prevented the extraction of features under unfortunate conditions, e. g., OPENFACE when no face appears or when only small faces appear in the original frame. Only the FASTTEXT features are repeated for the duration of a word and non-linguistic parts are also imputed with zero vectors. For the FASTTEXT alignment, the features are aggregated in such a way that for a FASTTEXT feature vector only one corresponding aggregated feature of any other type exists. This preparation should enable the participants to get started quickly and at the same time allows for own imputation procedures as well as unaligned modelling.

5 BASELINE SYSTEMS

For each Sub-challenge, a series of state-of-the-art approaches have been applied, and for reproducibility, all resources are made freely available⁶. In the proceeding section, we describe in detail the approaches. An overview of all baseline results is given in Table 2, Table 3, and Table 4. For both Sub-Challenges *MuSe-Wild*, and *MuSe-Trust*, the paradigm is continuous prediction of emotional signals. For this, we have applied a Recurrent Neural Network (RNN) with self-attention approach, and a deep audio-to-target end-to-end approach. In addition to these models, we use Support Vector

⁶<https://github.com/Istappen/MuSe2020>

Machines (SVMs), a multimodal Transformer and a fine-tuned NLP Transformer Albert to predict the classes of *MuSe-Topic*.

5.1 Early Fusion LSTM-RNN with Self-Attention

In order to address the sequential nature of the input features, we utilise a Long Short-Term Memory (LSTM)-RNN based architecture. The input feature sequences are input into two parallel LSTM-RNNs with hidden state dimensionality equal to 40, to encode the two corresponding query and value vector sequences. A self-attention sequence is calculated by means of a query and key dot product using a sequence-wide attention window. The attention and query sequences are then concatenated. For the continuous-time tasks *MuSe-Wild* and *MuSe-Trust*, the resulting hidden vector for each time step is further encoded by a feed-forward layer that outputs a one-dimensional prediction sequence per prediction target. For the *MuSe-Topic* task, we instead apply global max-pooling, to integrate the sequential information into one hidden state vector, which is then input into a feed-forward layer to provide the logits. In the former case, all the input samples are further segmented into 50 time-step sub-segments which are all used for training, whereas in the latter we pad/crop all sequences to 500 steps.

5.2 End-to-End Learning

As our end-to-end baseline we use End2You [45]: an open-source toolkit for multimodal profiling by end-to-end deep learning [43, 44]. For our purposes, we utilise three modalities, namely, audio, visual, and textual. Our audio model is inspired by a recently proposed emotion recognition model [46], and is comprised of a convolution recurrent neural network (CRNN). In particular, we use 3 convolution layers to extract spatial features from the raw segments. Our visual information is comprised of the VGGFACE features, where we use zero vectors when the face is not detected in a frame. Finally, as text features we use FASTTEXT, where we replicate the text features that span across several segments. We concatenate all uni-modal features and feed them to a one layer LSTM to capture the temporal dynamics in the data before the final prediction.

5.3 Multimodal Transformer

As baseline for the non-sequential predictions of *MuSe-Topic*, we choose the Multimodal Transformer (MMT) [42]. By using aligned and unaligned vision, language, and audio features for single label prediction, it outperformed state-of-the-art methods in a more text-focused Multimodal Sentiment Analysis setting. MMT merges multimodal timeseries using a feed-forward fusion process consisting of multiple crossmodal Transformer units. At the core of this network architecture are crossmodal attention modules which fuse multimodal features by directly attending to low-level features across all modalities. To predict topics, valence, and arousal we always utilise 3 feature sets, either of our three (tri), or of only two (bi) different modalities fed into the network. We noticed that after approximately 20 epochs the network converged. The model uses 5 crossmodal attention heads and an initial learning rate of 10^{-3} .

5.4 Albert

To reflect the current trend towards Transformer language models, such as Bidirectional Encoder Representations from Transformers

(BERT) [10], we include one of the latest versions, Albert [17], as a purely text-based baseline model. The authors of Albert proposed parameter reduction techniques, so that the total memory consumption is lower while increasing the training speed. These models supposedly scale better than the original BERT. The architecture is able to achieve state-of-the-art results on several benchmarks, despite having a relatively smaller number of parameters. For our purposes, we found a supervised tuning on the train partition for 3 epochs and balanced class weights to have the best effect. We applied a learning rate of 10^{-5} for the adjusted Adam Optimiser and set ϵ to 10^{-8} . With a sequence length of 300, the batch size has to be limited to 12 samples to be trained with 32GB GPU memory.

5.5 Support Vector Machines

For the task of emotion prediction in the Sub-Challenge *MuSe-Topic* only, we choose also to include results obtained through the use of conventional and easily reproducible Support Vector Machines (SVMs). These experiments employ the SCIKIT-LEARN toolkit, with a LINEARSVR classifier. No standardisation or normalisation was applied to any of the reported feature sets. The complexity parameter C was always optimised from 10^{-5} to 1 during the development phase, and the best value for C is reported. In contrast to our other approaches, we retrain the model on a concatenation of the train and development sets to predict the final test set result.

6 BASELINE RESULTS

6.1 MuSe-Wild

We evaluated several feature sets and combinations for the prediction of the continuous arousal and valence (see Table 2 for detailed results). For the prediction of arousal, the LSTM-RNN with self-attention using LLDs as input features, achieved the best result of all applied systems with a CCC of .3088 on the devel set, and .2884 on the test set. However, the combined metric (mean of valence and arousal) is considerably lower (CCC: .1931 on devel) due to the poor efficiency on the prediction of valence. Therefore, we define the **end-to-end framework utilising FASTTEXT, VGGFACE, and audio representations learnt from the raw audio signal as our baseline**, achieving a CCC of .2431 (on test) for the prediction of valence and .2706 (on test) for the prediction of arousal. We report a combined score of .2568 (on test) for this system.

6.2 MuSe-Topic

Table 3 shows the results of the baseline systems on the language-centric task of topic prediction. In line with recent research, the state-of-the-art NLP Transformer **Albert**, fine-tuned on the training set, achieved with 76.79% **(combined, on test) the best baseline result** leaving the second best system, the Multimodal Transformer utilising FASTTEXT, EGEMAPS, and FACIAL ACTION UNITS features (52.98% on test), far behind. The most successful configuration of the LSTM + Self-attention, the only not Transformer-based architecture, has another nearly 15% performance gap (on test 37.37%) to the MMT demonstrating the competitiveness of our baseline and the suitability of Transformers for this task.

For the task of emotion (valence and arousal) prediction in the *MuSe-Topic* sub-challenge, we also report baseline results in Table 4. Here, the picture is more balanced with some system failing to

Table 2: Reporting Arousal, Valence and Combined ($0.5 \cdot \text{Arousal} + 0.5 \cdot \text{Valence}$) for MuSe-Wild and Trustworthiness for MuSe-Trust, both using concordance correlation coefficient (CCC). As feature sets FASTTEXT (FT), EGEMAPS (Ge), DEEPSPECTRUM (DS), GoCAR (Go), VGGFACE (VG), and XCEPTION (X) and all visual (aV) are fed into the models. Furthermore, the raw audio signal (RA) is used in End2You, and LOW-LEVEL DESCRIPTORS (LDD) are utilised for MuSe-Trust in order to predict trustworthiness. All utilised features of MuSe-Wild and MuSe-Trust are aligned to the label timestamps by imputing missing values or repeating the word embeddings for FASTTEXT.

System	Features	Valence	Arousal	Combined	Trustworthiness
		devel / test	devel / test	devel / test	devel / test
LSTM + Self-ATT	LLD	.0711 / .0349	.3078 / .2834	.1894 / .1592	.2560 / .1343
LSTM + Self-ATT	DS	.0165 / .0024	.1585 / .1723	.0875 / .0874	.2019 / .1701
LSTM + Self-ATT	Ge	.0435 / -.0097	.1090 / .0827	.0762 / .0365	.1576 / .1385
LSTM + Self-ATT	FT	.1273 / .1816	.0959 / .1074	.1116 / .1445	.2278 / .2549
LSTM + Self-ATT	Ge + FT	.0520 / .0361	.1375 / .1018	.0947 / .0690	.2296 / .2054
LSTM + Self-ATT	X	.0499 / .0426	.0776 / .0683	.0638 / .0555	.1178 / .1664
LSTM + Self-ATT	aV	.0098 / .0272	.1598 / .1227	.0848 / .0749	.1167 / .1378
LSTM + Self-ATT	Ge + FT + V	.0393 / .0654	.1809 / .0865	.1101 / .0760	.1245 / .1695
End2You	FT + VG + RA	.1506 / .2431	.2587 / .2706	.2047 / .2568	.3198 / .4128
End2You-Multitask	FT + VG + RA	-	-	-	.3264 / .4119

Table 3: MuSe-Topic: Reporting Unweighted Average Recall (UAR), F1, and Combined ($0.66 \cdot F1 + 0.34 \cdot UAR$) for the topic predictions. As feature sets FASTTEXT (FT), Raw Text (RT), EGEMAPS (eG), DEEPSPECTRUM (DS), VGGFACE (VG), XCEPTION (X), OPENPOSE (OP), GoCAR (Go), FACIAL ACTION UNITS (AU) and all visual features (aV) are used. Two types of alignment are used to a) align to EGEMAPS (GA), and b) aggregate on FASTTEXT word features (FA).

System	Features	Alig.	F1	UAR	Combined
			devel / test	devel / test	devel / test
LSTM + Self-ATT	DS	GA	19.85 / 34.60	12.95 / 35.00	17.50 / 34.74
LSTM + Self-ATT	eG	GA	19.02 / 34.44	12.34 / 33.94	16.75 / 34.27
LSTM + Self-ATT	FT	GA	24.62 / 36.19	15.25 / 36.22	21.44 / 36.20
LSTM + Self-ATT	eG + FT	GA	20.38 / 35.32	13.13 / 34.87	17.92 / 35.16
LSTM + Self-ATT	X	GA	26.06 / 36.83	20.42 / 36.61	24.14 / 36.75
LSTM + Self-ATT	aV	GA	27.42 / 34.92	21.57 / 34.41	25.43 / 34.75
LSTM + Self-ATT	eG + FT + V	GA	27.58 / 37.14	20.08 / 37.14	25.03 / 37.14
Fine-tuned Albert	RT	-	71.69 / 76.59	69.56 / 77.18	70.96 / 76.79
MMT	FT + eG + X	-	48.24 / 53.18	41.49 / 50.44	44.86 / 51.81
MMT	FT + eG + X	FA	49.21 / 52.06	40.78 / 49.68	46.35 / 51.25
MMT	FT + eG + VG	-	44.72 / 50.71	35.60 / 45.19	41.62 / 48.84
MMT	FT + eG + Go	-	44.72 / 53.73	38.14 / 49.85	42.48 / 52.41
MMT	FT + eG + AU	-	46.22 / 54.52	40.66 / 49.99	44.33 / 52.98
MMT	FT + Go + OP	-	45.17 / 52.30	37.82 / 48.61	42.67 / 51.05
MMT	FT + DS + Go	-	44.42 / 52.06	36.77 / 49.59	41.82 / 51.22

achieve results higher than chance level (33 %) on test e. g., the fine-tuned Albert. Overall, the **Multimodal Transformer achieved with 38.81% (combined valence and arousal) the best results utilising FASTTEXT, EGEMAPS, and XCEPTION**. The same configuration is also **most successful in predicting valence (40.12%) on test**. The utilised SVMs, chosen due to their scalability on high dimensional data, showed results comparable to most state-of-the-art approaches. In particular, for the prediction of arousal, the **VG-GFACE features result is the best Combined F1 and UAR of 42.67 on test**. These SVM results lead us to assume that this task may benefit from a more traditional feature-level analysis. The confusion matrix for all tasks are depicted in Figure 2.

6.3 MuSe-Trust

The results for the prediction of trustworthiness are depicted in Table 2. Similar to MuSe-Wild, the end-to-end baseline system using FASTTEXT, VGGFACE, and raw audio signals gave the best results with .4128 CCC on test. The results may improve if the valence and arousal predicted signals are incorporated during training.

This can be accomplished in three ways: i) the model from MuSe-Wild is utilised to predict arousal and valence on MuSe-Trust; ii) the arousal and valence models can be retrained on MuSe-Trust (we provide train and devel labels); or iii) all three are predicted in a multitask-fashion (one model, 3 outputs) on train and devel, and only trustworthiness is predicted on test. We decided for option (iii). Adding these signals to the **end-to-end baseline system**, the predictive power of the model is similar to the previous one with CCC .3264 on the development set and .4119 on the test set.

7 CONCLUSIONS

In this paper, we introduced MuSe 2020 – the first Multimodal Sentiment Analysis in real media assessment challenge. MuSe 2020 utilises the MuSe-CaR multimodal corpus of emotional car reviews and comprises three Sub-challenges: i) MuSe-Wild, where the level of the affective dimensions of valence (corresponding to sentiment) and arousal has to be predicted from a ca. 35 hour data subset; ii) MuSe-Topic, where the domain-related conversational topic (10

Table 4: MuSe-Topic : Reporting Unweighted Average Recall (UAR), F1, and Combined ($0.66 \cdot F1 + 0.34 \cdot UAR$) for the 3-class valence and arousal predictions and the combined (mean) of valence and arousal. As feature sets FASTTEXT (FT), Raw Text (RT), EGEMAPS (eG), DEEPSPECTRUM (DS), VGGFACE (VG), XCEPTION (X), GoCAR (Go), OPENPOSE (OP), FACIAL ACTION UNITS (AU), and all visual feature set (aV) are used. Two types of alignment are used to a) align to EGEMAPS (GA) or b) aggregate on FASTTEXT word features (FA).

System	Features	Align.	c-Valence			c-Arousal			Combined
			F1	UAR	Combined	F1	UAR	Combined	
			devel / test	devel / test	devel / test	devel / test	devel / test	devel / test	
Fine-tuned Albert	RT		36.18 / 34.21	33.17 / 33.05	35.16 / 33.81	33.33 / 37.14	33.69 / 34.30	33.45 / 36.18	34.30 / 35.00
LSTM + Self-ATT	DS	GA	34.17 / 34.60	34.07 / 35.00	34.13 / 34.74	38.03 / 37.54	38.43 / 36.78	38.17 / 37.28	36.15 / 36.01
LSTM + Self-ATT	eG	GA	33.26 / 34.44	32.16 / 33.94	32.89 / 34.27	34.39 / 33.33	34.44 / 32.87	34.41 / 33.18	33.65 / 33.73
LSTM + Self-ATT	FT	GA	38.41 / 36.19	37.75 / 36.22	38.18 / 36.20	35.15 / 34.92	35.78 / 37.10	35.37 / 35.66	36.78 / 35.93
LSTM + Self-ATT	eG + FT	GA	34.92 / 35.32	34.05 / 34.87	34.63 / 35.16	34.39 / 35.48	34.48 / 35.42	34.42 / 35.46	34.53 / 35.31
LSTM + Self-ATT	X	GA	36.21 / 36.83	35.75 / 36.61	36.06 / 36.75	40.38 / 35.16	40.51 / 34.87	40.43 / 35.06	38.24 / 35.91
LSTM + Self-ATT	aV	GA	35.61 / 34.92	35.10 / 34.41	35.44 / 34.75	38.11 / 34.21	38.26 / 35.39	38.16 / 34.61	36.80 / 34.68
LSTM + Self-ATT	eG + FT + aV	GA	36.06 / 37.14	35.20 / 37.14	35.77 / 37.14	39.92 / 35.16	40.44 / 34.76	40.10 / 35.02	37.93 / 36.08
MMT	FT + eG + X		38.28 / 39.92	37.62 / 40.52	38.06 / 40.12	41.87 / 37.30	40.83 / 37.87	41.52 / 37.50	39.79 / 38.81
MMT	FT + eG + VG		37.38 / 32.78	38.19 / 32.53	37.65 / 32.69	47.12 / 41.19	45.55 / 39.01	46.58 / 40.45	42.12 / 36.57
MMT	DS + eG + VG		39.40 / 32.54	38.08 / 32.40	38.95 / 32.49	45.77 / 41.03	44.66 / 40.63	45.39 / 40.89	42.17 / 36.69
MMT	X + eG + VG		38.28 / 36.43	37.76 / 37.39	38.10 / 36.76	45.24 / 40.95	43.81 / 38.66	44.76 / 40.17	41.43 / 38.46
MMT	FT + eG + AU		36.93 / 39.92	37.35 / 39.57	37.07 / 39.80	43.15 / 34.76	41.88 / 34.87	42.72 / 34.80	39.89 / 37.30
MMT	FT + eG + OP		39.48 / 38.81	39.17 / 38.64	39.37 / 38.75	38.88 / 37.70	38.95 / 38.10	38.90 / 37.83	39.14 / 38.29
MMT	OP + eG + AU		37.30 / 36.67	36.34 / 37.45	36.97 / 36.93	43.15 / 34.68	42.01 / 35.69	42.76 / 35.03	39.87 / 35.98
End2You	FT + eG + X		37.19 / 33.54	35.70 / 33.18	36.68 / 33.42	42.76 / 32.45	42.67 / 33.34	42.73 / 32.75	39.70 / 33.08
SVM	eG		36.33 / 33.10	34.79 / 34.13	35.81 / 33.45	43.52 / 34.37	42.27 / 33.43	43.10 / 34.05	39.45 / 33.75
SVM	DS		34.08 / 34.29	33.21 / 34.07	33.79 / 34.21	41.35 / 42.30	40.18 / 40.18	40.18 / 41.83	36.98 / 38.02
SVM	X		38.28 / 37.94	37.09 / 37.94	37.87 / 37.94	46.22 / 41.35	45.25 / 40.52	45.89 / 41.07	41.88 / 39.50
SVM	VG		37.08 / 32.94	37.01 / 32.63	37.06 / 32.83	46.44 / 42.46	45.21 / 43.07	46.02 / 42.67	41.54 / 37.75
SVM	FT		37.90 / 36.43	36.00 / 35.37	37.26 / 36.07	45.17 / 38.25	44.53 / 39.67	44.95 / 38.74	41.10 / 37.40




Figure 2: Relative confusion matrix over all 10 topics of fine-tuned Albert (left) as well as the MMT (FASTTEXT, EGEMAPS and XCEPTION) for the prediction of valence (middle) and MMT (FASTTEXT, EGEMAPS and XCEPTION) for the prediction of arousal (right) classes on the test partition for sub-challenges MuSe-Topic.

classes) as well as three classes (low, medium and high) of valence and arousal have to be predicted from video parts containing the discussed topic; and, iii) MuSe-Trust, where the level of continuous trustworthiness has to be predicted from features and/or affective annotations. By intention, we decided to use open-source software to extract a wide range of feature sets to deliver the highest possible transparency and realism for the baselines. Besides the features, we also share the raw data and the developed code for our baselines on a public platform. Results indicate that: i) the level of affection in-the-wild is best predicted when the system is trained on the raw audio features; ii) for MuSe-Topic, (NLP-specific) Transformers are clearly superior when it comes to the prediction of topics, and no system is clearly outperforming on the three class valence and arousal prediction; and iii), in MuSe-Trust, adding valence and arousal contours as ‘signals’ in addition to other features is beneficial for the prediction of trustworthiness. The baselines also show the challenge ahead in mastering multimodal sentiment analysis, in particular when data are collected in user-generated, noisy environments. In the participants’ and future efforts, we expect novel exciting combinations of the modalities – potentially also such as linking modalities on earlier stages or more closely.

8 ACKNOWLEDGMENTS

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 115902 (RADAR CNS) and No. 826506 (sustAGE), the EP-SRC Grant No. 2021037, and the Bavarian State Ministry of Education, Science and the Arts in the framework of the Centre Digitisation.Bavaria (ZD.B). We thank the sponsors of the Challenge BMW Group and audeERING.

REFERENCES

- [1] Luis Aguado, Francisco J Román, María Fernández-Cahill, Teresa Diéguez-Risco, and Verónica Romero-Ferreiro. 2011. Learning about faces: effects of trustworthiness on affective evaluation. *The Spanish journal of psychology* 14, 2 (2011), 523–534.
- [2] Shahin Amiriparian, Maurice Gerczuk, Sandra Ottl, Nicholas Cummins, Michael Freitag, Sergey Pugachevskiy, Alice Baird, and Björn W Schuller. 2017. Snore Sound Classification Using Image-Based Deep Spectrum Features.. In *INTER-SPEECH*, Vol. 434. 3512–3516.
- [3] John Arevalo, Thamar Solorio, Manuel Montes-y Gómez, and Fabio A González. 2020. Gated multimodal networks. *Neural Computing and Applications* (2020), 1–20.
- [4] T. Baltrušaitis, P. Robinson, and L. Morency. 2016. OpenFace: an Open Source Facial Behavior Analysis Toolkit. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*. IEEE, Lake Placid, NY. 10 pages.
- [5] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics* 5 (2017), 135–146.
- [6] Qiong Cao, Li Shen, Weidi Xie, Omkar M. Parkhi, and Andrew Zisserman. 2017. VGGFace2: A dataset for recognising faces across pose and age. *CoRR abs/1710.08092* (2017). arXiv:1710.08092 <http://arxiv.org/abs/1710.08092>
- [7] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. 2019. OpenPose: Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2019).
- [8] Nicholas Cummins, Shahin Amiriparian, Gerhard Hagerer, Anton Batliner, Stefan Steidl, and Björn W Schuller. 2017. An image-based deep spectrum feature representation for the recognition of emotional speech. In *Proceedings of the 25th ACM international conference on Multimedia*. 478–484.
- [9] Maria Teresa Cuomo, Debora Tortora, Alex Giordano, Giuseppe Festa, Gerardino Metallo, and Erika Martinelli. 2020. User-generated content in the era of digital well-being: A netnographic analysis in a healthcare marketing context. *Psychology & Marketing* (2020).
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- [11] Florian Eyben, Klaus R Scherer, Björn W Schuller, Johan Sundberg, Elisabeth André, Carlos Busso, Laurence Y Devillers, Julien Epps, Petri Laukka, Shrikanth S Narayanan, et al. 2015. The Geneva minimalistic acoustic parameter set (GeMAPS) for voice research and affective computing. *IEEE transactions on affective computing* 7, 2 (2015), 190–202.
- [12] Florian Eyben, Martin Wöllmer, and Björn Schuller. 2010. Opensmile: the munich versatile and fast open-source audio feature extractor. In *Proceedings of the 18th ACM international conference on Multimedia*. 1459–1462.
- [13] Raul Gomez, Jaume Gibert, Lluís Gomez, and Dimosthenis Karatzas. 2020. Exploring Hate Speech Detection in Multimodal Publications. In *The IEEE Winter Conference on Applications of Computer Vision*. 1470–1478.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [15] Dimitrios Kollias, Attila Schul, Elnar Hajiyev, and Stefanos Zafeiriou. 2020. Analysing affective behavior in the first ABAW 2020 competition. *arXiv preprint arXiv:2001.11409* (2020).
- [16] Dimitrios Kollias, Panagiotis Tzirakis, Mihalís 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* (2019), 1–23.
- [17] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=H1eA7AEtVS>
- [18] I Lawrence and Kuei Lin. 1989. A concordance correlation coefficient to evaluate reproducibility. *Biometrics* (1989), 255–268.
- [19] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*. Springer, 740–755.
- [20] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. 2015. Deep Learning Face Attributes in the Wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- [21] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.
- [22] Saif M Mohammad. 2016. Sentiment analysis: Detecting valence, emotions, and other affectual states from text. In *Emotion measurement*. Elsevier, 201–237.
- [23] Vedhas Pandit and Björn Schuller. 2019. On Many-to-Many Mapping Between Concordance Correlation Coefficient and Mean Square Error. *arXiv preprint arXiv:1902.05180* (2019).
- [24] Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman. 2015. Deep Face Recognition. In *Proceedings of the British Machine Vision Conference (BMVC)*, Mark W. Jones Xianghua Xie and Gary K. L. Tam (Eds.). BMVA Press, Article 41, 12 pages. <https://doi.org/10.5244/C.29.41>
- [25] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.
- [26] Frits K Pil and Matthias Holweg. 2004. Linking product variety to order-fulfillment strategies. *Interfaces* 34, 5 (2004), 394–403.
- [27] Daniel Preotjuc-Pietro, H Andrew Schwartz, Gregory Park, Johannes Eichstaedt, Margaret Kern, Lyle Ungar, and Elisabeth Shulman. 2016. Modelling valence and arousal in facebook posts. In *Proceedings of the 7th workshop on computational approaches to subjectivity, sentiment and social media analysis*. 9–15.
- [28] Xiaoyu Qiu, Zhiqun Feng, Xiaohui Yang, and Jinglan Tian. 2020. Multimodal Fusion of Speech and Gesture Recognition based on Deep Learning. In *Journal of Physics: Conference Series*, Vol. 1453. 012092.
- [29] Joseph Redmon and Ali Farhadi. 2018. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767* (2018).
- [30] Fabien Ringeval, Björn Schuller, Michel Valstar, Jonathan Gratch, Roddy Cowie, Stefan Scherer, Sharon Mozgai, Nicholas Cummins, Maximilian Schmitt, and Maja Pantic. 2017. Avec 2017: Real-life depression, and affect recognition workshop and challenge. In *Proceedings of the 7th Annual Workshop on Audio/Visual Emotion Challenge*. 3–9.
- [31] Fabien Ringeval, Andreas Sonderegger, Juergen Sauer, and Denis Lalanne. 2013. Introducing the RECOLA multimodal corpus of remote collaborative and affective interactions. In *2013 10th IEEE international conference and workshops on automatic face and gesture recognition (FG)*. IEEE, 1–8.
- [32] James A Russell. 1980. A circumplex model of affect. *Journal of personality and social psychology* 39, 6 (1980), 1161.
- [33] Björn Schuller, Stefan Steidl, Anton Batliner, Alessandro Vinciarelli, Klaus Scherer, Fabien Ringeval, Mohamed Chetouani, Felix Weninger, Florian Eyben, Erik Marchi, et al. 2013. The INTERSPEECH 2013 computational paralinguistics challenge: Social signals, conflict, emotion, autism. In *Proceedings INTERSPEECH 2013, 14th Annual Conference of the International Speech Communication Association, Lyon, France*.
- [34] Björn W Schuller. 2013. *Intelligent audio analysis*. Springer.
- [35] Björn W Schuller, Anton Batliner, Christian Bergler, Eva-Maria Messner, Antonia Hamilton, Shahin Amiriparian, Alice Baird, Georgios Rizos, Maximilian Schmitt, Lukas Stappen, et al. 2020. The INTERSPEECH 2020 Computational Paralinguistics Challenge: Elderly Emotion, Breathing & Masks. *Proceedings INTERSPEECH Shanghai, China: ISCA* (2020).
- [36] Björn W Schuller, Stefan Steidl, Anton Batliner, Peter B Marschik, Harald Baumeister, Fengquan Dong, Simone Hantke, Florian B Pokorny, Eva-Maria Rathner, Katrin D Bartl-Pokorny, et al. 2018. The INTERSPEECH 2018 Computational Paralinguistics Challenge: Atypical & Self-Assessed Affect, Crying & Heart Beats.. In *Interspeech*. 122–126.
- [37] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [38] Mohammad Soleymani, David Garcia, Brendan Jou, Björn Schuller, Shih-Fu Chang, and Maja Pantic. 2017. A survey of multimodal sentiment analysis. *Image and Vision Computing* 65 (2017), 3–14.
- [39] Lukas Stappen, Xinchen Du, Vincent Karas, Stefan Müller, and Björn W Schuller. 2020. Go-CaRD—Generic, Optical Car Part Recognition and Detection: Collection, Insights, and Applications. *arXiv preprint arXiv:2006.08521* (2020).
- [40] Lukas Stappen, Vincent Karas, Nicholas Cummins, Fabien Ringeval, Klaus Scherer, and Björn Schuller. 2019. From speech to facial activity: towards cross-modal sequence-to-sequence attention networks. In *2019 IEEE 21st International Workshop on Multimedia Signal Processing (MMSp)*. IEEE, 1–6.
- [41] Mike Thelwall, Kevan Buckley, Georgios Paltoglou, Di Cai, and Arvid Kappas. 2010. Sentiment strength detection in short informal text. *Journal of the American society for information science and technology* 61, 12 (2010), 2544–2558.
- [42] Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J. Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019. Multimodal Transformer for Unaligned Multimodal Language Sequences. *CoRR abs/1906.00295* (2019). arXiv:1906.00295 <http://arxiv.org/abs/1906.00295>
- [43] Panagiotis Tzirakis, George Trigeorgis, Mihalís A Nicolaou, Björn W Schuller, and Stefanos Zafeiriou. 2017. End-to-end multimodal emotion recognition using deep neural networks. *IEEE Journal of Selected Topics in Signal Processing* 11, 8 (2017), 1301–1309.
- [44] Panagiotis Tzirakis, Stefanos Zafeiriou, and Björn Schuller. 2019. Real-world automatic continuous affect recognition from audiovisual signals. In *Multimodal Behavior Analysis in the Wild*. Elsevier, 387–406.
- [45] Panagiotis Tzirakis, Stefanos Zafeiriou, and Björn W Schuller. 2018. End2You—The Imperial Toolkit for Multimodal Profiling by End-to-End Learning. *arXiv preprint arXiv:1802.01115* (2018).

- [46] Panagiotis Tzirakis, Jiehao Zhang, and Bjorn W Schuller. 2018. End-to-end speech emotion recognition using deep neural networks. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 5089–5093.
- [47] Michel Valstar, Björn Schuller, Kirsty Smith, Florian Eyben, Bihan Jiang, Sanjay Bilakhia, Sebastian Schnieder, Roddy Cowie, and Maja Pantic. 2013. AVEC 2013: the continuous audio/visual emotion and depression recognition challenge. In *Proceedings of the 3rd ACM international workshop on Audio/visual emotion challenge*. ACM, 3–10.
- [48] Zheng Wang, Jie Zhou, Jing Ma, Jingjing Li, Jiangbo Ai, and Yang Yang. 2020. Discovering attractive segments in the user-generated video streams. *Information Processing & Management* 57, 1 (2020), 102130.
- [49] Shuo Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. 2015. WIDER FACE: A Face Detection Benchmark. *CoRR* abs/1511.06523 (2015). arXiv:1511.06523 <http://arxiv.org/abs/1511.06523>
- [50] Amir Zadeh, Paul Pu Liang, Louis-Philippe Morency, Soujanya Poria, Erik Cambria, and Stefan Scherer. 2018. Proceedings of Grand Challenge and Workshop on Human Multimodal Language (Challenge-HML). In *Proceedings of Grand Challenge and Workshop on Human Multimodal Language (Challenge-HML)*.
- [51] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. 2016. Joint Face Detection and Alignment Using Multitask Cascaded Convolutional Networks. *IEEE Signal Processing Letters* 23 (04 2016). <https://doi.org/10.1109/LSP.2016.2603342>