

SPEECH EMOTION RECOGNITION USING SEMANTIC INFORMATION

Panagiotis Tzirakis¹, Anh Nguyen¹, Stefanos Zafeiriou¹, Björn W. Schuller^{1,2}

¹ GLAM – Group on Language, Audio, & Music, Imperial College London, UK

² EIHW – Chair of Embedded Intelligence for Health Care and Wellbeing, University of Augsburg, Germany
email: panagiotis.tzirakis12@imperial.ac.uk

ABSTRACT

Speech emotion recognition is a crucial problem manifesting in a multitude of applications such as human computer interaction and education. Although several advancements have been made in the recent years, especially with the advent of Deep Neural Networks (DNN), most of the studies in the literature fail to consider the semantic information in the speech signal. In this paper, we propose a novel framework that can capture both the semantic and the paralinguistic information in the signal. In particular, our framework is comprised of a semantic feature extractor, that captures the semantic information, and a paralinguistic feature extractor, that captures the paralinguistic information. Both semantic and paralinguistic features are then combined to a unified representation using a novel attention mechanism. The unified feature vector is passed through a LSTM to capture the temporal dynamics in the signal, before the final prediction. To validate the effectiveness of our framework, we use the popular SEWA dataset of the AVEC challenge series and compare with the three winning papers. Our model provides state-of-the-art results in the valence and liking dimensions.¹

Index Terms— emotion recognition, deep learning, semantic, paralinguistic, audiotextual information

1. INTRODUCTION

Automatic affect recognition is a vital component in human-to-human communication affecting our social interaction, perception among others [1]. In order to accomplish a *natural* interaction between human and machine, intelligent systems need to recognise the emotional state of individuals. However, the task is challenging, as human emotions lack of temporal boundaries and different individuals express emotions in different ways [2]. In addition, emotions are expressed through multiple modalities. Over the past two decades, a plethora of systems have been proposed that utilise several modalities such as physiological signals, facial expression, speech, and text [3, 4, 5, 2, 6]. To achieve an accurate emotion recognition system, it is important to consider multiple modalities, as complementary information exists among them [3].

Current studies exploit Deep Neural Networks (DNNs) to model affect using multiple modalities [7, 8, 9, 10]. Two modalities that have been extensively used for the emotion recognition task are speech and text [11, 12]. Whereas the speech signal provides low-level characteristics of the emotions (e. g., prosody), text provides high-level (semantic) information (e. g., the words “love” and “like” carry strong emotional content). To this end, several systems have

shown that integrating both modalities, strong performance gains can be obtained [11].

However, one may argue that the textual information is redundant, as it is already included in the speech signal, and as such semantic information can be captured using only the speech modality. To this end, we propose an audiotextual training framework, where the text modality is used during training, but discarded during evaluation. In particular, we train Word2Vec [13] and Speech2Vec [14] models, and align their two embedding spaces such that Speech2Vec features are as close as possible with the Word2Vec ones [15]. In addition to the semantic information, we capture low-level characteristics of the speech signal by training a convolution recurrent neural network. The semantic and paralinguistic features are combined to a unified representation and passed through a long short-term memory (LSTM) module that captures the temporal dynamics in the signal, before the final prediction.

To test the effectiveness of our model, we utilise the Sentiment Analysis in the Wild (SEWA) dataset, which was used in the Audio/Visual Emotion Challenge (AVEC) since 2017 [16]. The dataset provides three continuous affect dimensions: arousal, valence, and likability. Although the arousal and valence dimensions are easily integrated in a single network during the training phase of the models, the likability dimension can cause convergence and generalisation difficulties [16, 17]. To this end, we propose to use a novel ‘disentangled’ attention mechanism to fuse the semantic and paralinguistic features such that the information required per affect dimension is *disentangled*. Our approach provides training stability, and, at the same time, increases the generalisability of the network during evaluation. We compare our framework with the three best performing papers of the competition [18, 19, 17] in terms of concordance correlation coefficient (ρ_c) [20, 21], and show that our method provides state-of-the-art results for the valence and likability dimensions.

In summary, the main contributions of the paper are the following: (a) propose to use the acoustic speech signal to capture semantic information that exists in the text modality, (b) show how to disentangle the information in the network per affect dimensions for stable training and generalisability during the evaluation phase, and (c) produce state-of-the-art results in the valence and likability dimensions using the SEWA dataset.

2. RELATED WORK

Several studies have been proposed in the literature for speech emotion recognition [22, 23, 24]. For example, Trigeorgis et al. [22] utilised a convolution neural network to capture the spatial information in the signal, and a recurrent neural network for the temporal

¹Code available here: https://github.com/glam-imperial/semantic_speech_emotion_recognition

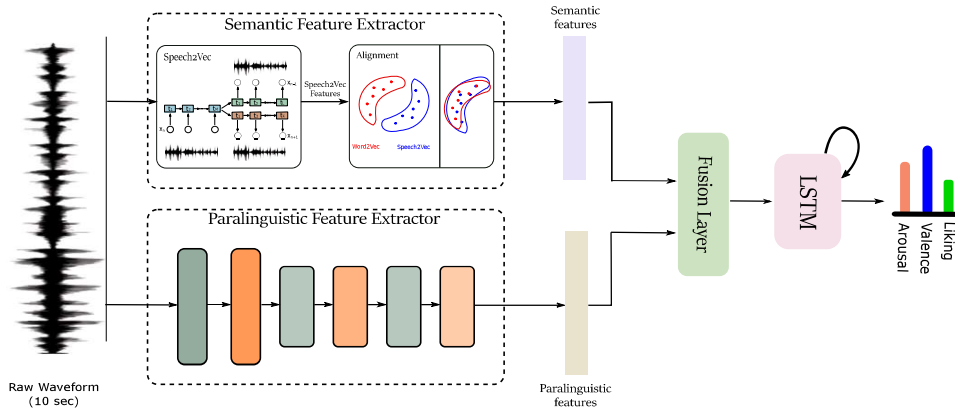


Fig. 1. Our proposed model is comprised of two networks: (a) the semantic feature extractor, that extracts high-level features containing semantic information of the input, and (b) the paralinguistic feature extractor, that extracts low-level features containing paralinguistic information of the signal. Both feature vectors are passed through a fusion layer, that combines the information and extracts a unified representation of the input, that is then passed through a LSTM model for the final prediction.

ones. In a similar study, Tzirakis et al. [23] showed that utilising a deeper architecture with longer input window produces better results. In another study, Neumann et al. [25] propose an attentive convolutional neural network (ACNN) that combines CNNs with attention.

In the past ten years, a plethora of models have been proposed that incorporate more than one modality for the emotion recognition task [8, 26, 9]. In particular, Tzirakis et al. [8] uses both audio and visual information for continuous emotion recognition. Although this study produced good results, it utilises both modalities for the training and evaluation of the model. In a more recent study, Albanie et al. [9] transfer the knowledge from the visual information (facial expressions) to the speech model. In another study, Han et al. [26] proposed an implicit fusion strategy for audiovisual emotion recognition. In this study, both audio and visual modalities are used for the training of the model and only one for the evaluation of the model.

3. PROPOSED METHOD

Our cross-modal framework can leverage the semantic (high-level) information (Sec. 3.1) and the paralinguistic (low-level) dynamics in the speech signal (Sec. 3.2). The low- and high-level feature sets are fused together using a novel attention fusion strategy (Sec. 3.3) before feeding them to a one-layer LSTM module, to capture the temporal dynamics in the signal, for the final frame-level prediction. Fig. 1 depicts the proposed method.

3.1. Semantic Feature Extractor

To capture the semantic information in the speech signal, we train Word2Vec and Speech2Vec models. The first model uses the text information to extract a semantic vector representation from a given word, whereas the second one the speech. We align their embedding spaces, similar to [15], for semantically richer speech representations. Mathematically, we define the speech embedding matrix $S = [s_1, s_2, \dots, s_m] \in \mathbb{R}^{m \times d_s}$ to be of m vocabulary words with dimension d_s , and the text embedding matrix $T = [t_1, t_2, \dots, t_n] \in \mathbb{R}^{n \times d_t}$

to be of n vocabulary words with dimension d_t . Our goal is to learn a linear mapping $W \in \mathbb{R}^{d_t \times d_s}$ such that WS is most similar to T .

To this end, we learn an initial proxy of W via domain-adversarial training. The adversarial training is a two-layer game where the generator tries, by computing W , to deceive the discriminator from correctly identifying the embedding space, and making WS and T as similar as possible. Mathematically, the discriminator tries to minimise the following objective:

$$L_D(\theta_D|W) = -\frac{1}{n} \sum_{i=1}^n \log P_{\theta_D}(\text{speech} = 1|Ws_i) - \frac{1}{m} \sum_{i=1}^m \log P_{\theta_D}(\text{speech} = 0|t_i), \quad (1)$$

where θ_D are the parameters of the discriminator, and $P_{\theta_D}(\text{speech} = 1|z)$ is the probability the vector z originates from speech embedding.

On the other hand, the generator tries to minimise the following objective:

$$L_G(W|\theta_D) = -\frac{1}{n} \sum_{i=1}^n \log P_{\theta_D}(\text{speech} = 0|Ws_i) - \frac{1}{m} \sum_{i=1}^m \log P_{\theta_D}(\text{speech} = 1|t_i). \quad (2)$$

A limitation of the above formulation is that all embedding vectors are treated equally during training. However, words with higher frequency would have better embedding quality in the vector space than less frequent words. To this end, we use the frequent words to create a dictionary that specifies which speech embedding vectors correspond to which text embedding vectors, and refine W :

$$W^* = \underset{W}{\operatorname{argmin}} \|WS_r - T_r\|_F, \quad (3)$$

where S_r is a matrix built by k speech vectors from S and T_r is a matrix built by k vectors from T . The solution of Eq. 3 is obtained from

Layer	Kernel/Stride	Channels	Activation
Convolution	8/1	50	ReLU
Max-pooling	10/10	—	—
Convolution	6/1	125	ReLU
Max-pooling	5/5	—	—
Convolution	6/1	125	ReLU
Max-pooling	5/5	—	—

Table 1. Paralinguistic feature extractor. Shown are the layer type, kernel/stride size, channels size, and activation function.

the singular value decomposition of $S_r T_r^T$, i. e., $SVD(S_r T_r^T) = U \Sigma V^T$.

3.2. Paralinguistic Feature Extractor

Our paralinguistic feature extraction network is comprised of three 1-D CNN layers with a rectified linear unit (ReLU) as activation function, and max-pooling operations in-between. Both convolution and pooling operations are performed on the time domain, using the raw waveform as input. Inspired by our previous work [23], we perform convolution with small kernel size and stride of one, and a large kernel and stride size for the max-pooling. Table 1 shows the architecture of the network.

3.3. Fusion Strategies

Our last step is to fuse the semantic ($\mathbf{x}_s \in R^{d_s}$) and paralinguistic ($\mathbf{x}_p \in \mathbb{R}^{d_p}$) speech features, before feeding them to the LSTM. This is performed with two strategies: (i) concatenation, (ii) ‘disentangled’ attention mechanism.

Concatenation. The first approach is a standard feature-level fusion, i. e., a simple concatenation of the feature vectors. Mathematically, $\mathbf{x}_{fusion} = [\mathbf{x}_s, \mathbf{x}_p]$.

Disentangled attention mechanism. For our second approach, we propose using attention mechanism to fuse the two modalities. To this end, we perform a linear projection for each of the feature sets such that they are in the same vector space (with dimension d_u):

$$\begin{aligned}\tilde{\mathbf{x}}_s &= \mathbf{W}_s \mathbf{x}_s + b_s, \\ \tilde{\mathbf{x}}_p &= \mathbf{W}_p \mathbf{x}_p + b_p,\end{aligned}\quad (4)$$

where $\mathbf{W}_s \in R^{d_u \times d_s}$, $\mathbf{W}_p \in R^{d_u \times d_p}$ are projection matrices for the semantic and paralinguistic feature sets, respectively.

We fuse these features using attention mechanism, i. e.,

$$\begin{aligned}Attention(\tilde{\mathbf{x}}_s, \tilde{\mathbf{x}}_t) &= \alpha_s \tilde{\mathbf{x}}_s + \alpha_p \tilde{\mathbf{x}}_p, \\ \alpha_i &= softmax\left(\frac{\tilde{\mathbf{x}}_i \mathbf{q}_i}{\sqrt{d_u}}\right),\end{aligned}\quad (5)$$

where $\mathbf{q} \in R^{d_u}$ is learnable vector that attends to different features.

At this point, we use three fully-connected (FC) layers with linear activation of same dimensionality on top of the output obtained from first attention layer, i. e.,

$$\begin{aligned}\mathbf{a} &= \mathbf{W}_a \tilde{\mathbf{x}}_{sp} + b_a \\ \mathbf{v} &= \mathbf{W}_v \tilde{\mathbf{x}}_{sp} + b_v, \\ \mathbf{l} &= \mathbf{W}_l \tilde{\mathbf{x}}_{sp} + b_l,\end{aligned}\quad (6)$$

where $\{\mathbf{W}_a, \mathbf{W}_l, \mathbf{W}_v\} \in R^{d_u \times d_u}$ are projection matrices.

We choose to use three FC layers such that the information flow per emotional dimension (i. e., arousal, valence, and liking) in the network is disentangled. The intuition here is that by adding three additional dense layers, we hope that each of these projections could learn features that suit best for a dimension in our emotion space. In case of a higher number of outputs, more FC layers can be used.

To fuse the information of the ‘disentangled’ vector spaces, we apply an attention layer so that each suited feature set could attend to one another and produce an enriched fusion feature output for final prediction. In particular, we, first, apply attention on \mathbf{a} and \mathbf{l} ; and, finally, on the result with \mathbf{v} , i. e.,

$$\begin{aligned}\mathbf{z} &= Attention(\mathbf{a}, \mathbf{l}) \\ \mathbf{x}_{fusion} &= Attention(\mathbf{z}, \mathbf{v}).\end{aligned}\quad (7)$$

4. DATASET

We test the performance of our proposed framework on a time-continuous emotion recognition dataset for real-world environments. In particular, as outlined, we utilise the Sentiment Analysis in the Wild (SEWA) dataset that was used in the AVEC 2017 challenge [16]. The dataset consists of ‘in-th-wild’ audiovisual recordings that were captured from web-cameras and microphones from 32 pairs (i. e., 64 participants) that watched a 90 sec commercial visual and discussed it with their partner for maximum of 3 min. It provides three modalities, namely, audio, visual, and text, for three emotional dimensions: arousal, valence, and liking. The dataset is split into 3 partitions: training (17 pairs), development (7 pairs), and test (8 pairs), and was annotated by 6 German-speaking annotators (3 female, 3 male).

5. EXPERIMENTS

5.1. Experimental Setup

For training the models, we utilised the Adam optimisation method [27], and a fixed learning rate of 10^{-4} throughout all experiments. We used a mini-batch of 25 samples with sequence length of 300, and a dropout [28] with $p = 0.5$ for all layers except the recurrent ones to regularise our network. This step is important, as our models have a large amount of parameters and not regularising the network makes it prone on overfitting on the training data. In addition, the LSTM network we use in the training phase is trained with a dropout of 0.5 and a gradient norm clipping of 5.0. Finally, we segment the raw waveform into 10 sec long sequences with sampling rate of 22 050 Hz. Hence, each sequence corresponds to a 22 0500-dimension vector.

5.2. Objective Function

Our objective function is based on the Concordance Correlation Coefficient (ρ_c) that was also used in the AVEC 2017 challenge. ρ_c evaluates the agreement level between the predictions and the gold standard by scaling their correlation coefficient with their mean square difference. Mathematically, the the concordance loss J_c can be defined as follows:

$$\mathcal{L}_c = 1 - \rho_c = 1 - \frac{2\sigma_{xy}^2}{\sigma_x^2 + \sigma_y^2 + (\mu_x - \mu_y)^2}, \quad (8)$$

where $\mu_x = \mathbb{E}(\mathbf{x})$, $\mu_y = \mathbb{E}(\mathbf{y})$, $\sigma_x^2 = \text{var}(\mathbf{x})$, $\sigma_y^2 = \text{var}(\mathbf{y})$, and $\sigma_{xy}^2 = \text{cov}(\mathbf{x}, \mathbf{y})$.

Our end-to-end network is trained to predict the arousal, valence, and liking dimensions, and as such, we define the overall loss as follows, $\mathcal{L} = (\mathcal{L}_c^a + \mathcal{L}_c^v + \mathcal{L}_c^l)/3$, where \mathcal{L}_c^a , \mathcal{L}_c^v , and \mathcal{L}_c^l are the concordance loss of the arousal, valence, and liking dimensions, respectively, contributing equally to the loss.

5.3. Ablation Study

5.3.1. Comparing Vector Spaces

We test the performance of both the semantic and paralinguistic networks, trained independently, and trained jointly, to show the beneficial properties of our proposed framework. Table 2 depicts the results in terms of ρ_c on the development set of the SEWA dataset. We observe that Word2Vec produces slightly better results than Speech2Vec. However, after aligning their embedding spaces, the aligned Speech2Vec has higher performance than Word2Vec, indicating both that the refinement process makes speech embedding similar to the word ones, and that paralinguistic information exists in the model. Finally, the paralinguistic network, although it produces worse results than the aligned Speech2Vec model, provides the best results for the arousal dimension.

Model	Arousal	Valence	Liking	Avg
Word2Vec	.434	.513	.208	.385
Speech2Vec	.433	.470	.182	.362
Align Speech2Vec	.453	.452	.257	.387
Paralinguistic	.508	.436	.154	.366

Table 2. SEWA dataset results (in terms of ρ_c) of the Word2Vec, Speech2Vec, aligned Speech2Vec and paralinguistic models, on the development set.

5.3.2. Fusion Strategies

We further explore the effectiveness of the attention fusion strategy compared to the simple concatenation. For our experiments we utilised both semantic and paralinguistic deep network models of the proposed method. Table 5.3.2 depicts the results in terms of ρ_c on the development set of the SEWA dataset. We observe that the attention method performs superior to the other one on all emotional dimensions, indicating the effectiveness of our approach to model the three emotional dimensions by projecting them to three different spaces before fusing them together with attention.

Fusion strategy	Arousal	Valence	Liking	Avg
Concatenation	.427	.428	.306	.387
Disentangled attention	.499	.497	.311	.435

Table 3. SEWA dataset results (in terms of ρ_c) of the various fusion methods (i. e., concatenation, attention, and hierarchical attention) in the development set.

Method	Arousal	Valence	Liking
Baseline [16]	.225 (.344)	.224 (.351)	-.020 (.081)
Dang et. al. [18]	.344 (.494)	.346 (.507)	— (—)
Huang et. al. [19]	.583 (.584)	.487 (.585)	— (—)
Chen et. al. [17]	.422 (.524)	.405 (.504)	.054 (.273)
Proposed	.429 (.499)	.503 (.497)	.312 (.311)

Table 4. SEWA dataset test result (in terms of ρ_c) of our proposed fusion model compared with the winning models in AVEC 2017. In parenthesis are the performances obtained on the development set. A dash is inserted if the results could not be obtained.

5.4. Results

We compare our proposed framework with the winning papers of the AVEC 2017 challenge. As our model utilises only the audio modality during evaluation, we show, for fairness of comparison, the results of these studies using the audio information. Table 4 depicts the results. First, we observe that our approach provides the best results in the valence dimension with high margin, and the second best in the arousal one. We should note, however, that the network from Huang et al. [19] was pretrained on 300 hours of a spontaneous English speech recognition corpus before fine-tuning it to the SEWA dataset. In addition to the features of the network, they also utilise several hand-engineered features. Second, we observe that our approach provides the highest performance in the likability dimension. Our method is able to generalise on this dimension compared to Chen et al [17] whose performance drops significantly compared to its performance on the development set. Finally, we should note the high generalisation capability of our approach to model all three emotional dimensions, indicating the effectiveness of the proposed disentangled attention mechanism strategy.

6. CONCLUSIONS

In this paper, we propose a training framework using audio and text information for speech emotion recognition. In particular, we use Word2Vec and Speech2Vec models, and align their embedding spaces for accurate semantic feature extraction using only the speech signal. We combine the semantic and paralinguistic features using a novel attention fusion strategy that first disentangles the information per emotional dimension, and then combines it using attention. The proposed model is evaluated on the SEWA dataset and produces state-of-the-art results on the valence and liking dimensions, when compared with the best performing papers submitted to the AVEC 2017 challenge.

In future work, we intend to use a single network to simultaneously capture the semantic and the paralinguistic information in the speech signal. This will result in simplifying, and at the same time, reducing the number of parameters of the model. Additionally, we intend to investigate the performance of the proposed method on categorical emotion recognition datasets.

7. ACKNOWLEDGEMENTS

The support of the EPSRC Center for Doctoral Training in High Performance Embedded and Distributed Systems (HiPEDS, Grant Reference EP/L016796/1) is gratefully acknowledged.

REFERENCES

- [1] R. Picard, *Affective Computing*, MIT Press, 1997.
- [2] C.-N. Anagnostopoulos, T. Iliou, and I. Giannoukos, “Features and classifiers for emotion recognition from speech: a survey from 2000 to 2011,” *Artificial Intelligence Review*, pp. 155–177, 2015.
- [3] P. Tzirakis, S. Zafeiriou, and B. Schuller, “Chapter 18 - real-world automatic continuous affect recognition from audiovisual signals,” in *Multimodal Behavior Analysis in the Wild*, pp. 387–406. 2019.
- [4] D. Kollias, P. Tzirakis, Mihalis A Nicolaou, A. Papaioannou, G. Zhao, B. Schuller, I. Kotsia, and S. Zafeiriou, “Deep affect prediction in-the-wild: Aff-wild database and challenge, deep architectures, and beyond,” *International Journal of Computer Vision*, pp. 907–929, 2019.
- [5] B. Schuller, “Speech emotion recognition: two decades in a nutshell, benchmarks, and ongoing trends,” *Communications of the ACM*, pp. 90–99, 2018.
- [6] J. Zhang, Z. Yin, P. Chen, and S. Nichele, “Emotion recognition using multi-modal data and machine learning techniques: A tutorial and review,” *Information Fusion*, pp. 103–126, 2020.
- [7] J. Han, Z. Zhang, M. Schmitt, M. Pantic, and B. Schuller, “From hard to soft: Towards more human-like emotion recognition by modelling the perception uncertainty,” in *Proc. ACM Multimedia*, 2017, pp. 890–897.
- [8] P. Tzirakis, G. Trigeorgis, M. Nicolaou, B. Schuller, and S. Zafeiriou, “End-to-end multimodal emotion recognition using deep neural networks,” *IEEE Journal of Selected Topics in Signal Processing*, pp. 1301–1309, 2017.
- [9] S. Albanie, A. Nagrani, A. Vedaldi, and A. Zisserman, “Emotion recognition in speech using cross-modal transfer in the wild,” in *Proc. ACM Multimedia*, 2018, pp. 292–301.
- [10] P. Tzirakis, S. Zafeiriou, and B. Schuller, “End2You—The Imperial Toolkit for Multimodal Profiling by End-to-End Learning,” *arXiv preprint arXiv:1802.01115*, 2018.
- [11] S. Yoon, S. Byun, and K. Jung, “Multimodal speech emotion recognition using audio and text,” in *Proc. SLT*, 2018, pp. 112–118.
- [12] M. V. Mäntylä, D. Graziotin, and M. Kuutila, “The evolution of sentiment analysis—A review of research topics, venues, and top cited papers,” *Computer Science Review*, pp. 16–32, 2018.
- [13] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, “Distributed representations of words and phrases and their compositionality,” in *Proc. Advances in neural information processing systems (NeurIPS)*, 2013, pp. 3111–3119.
- [14] Y.-A. Chung and J. Glass, “Speech2vec: A sequence-to-sequence framework for learning word embeddings from speech,” *Proc. Interspeech*, pp. 811–815, 2018.
- [15] Y.-A. Chung, W.-H. Weng, S. Tong, and J. Glass, “Unsupervised cross-modal alignment of speech and text embedding spaces,” in *Proc. Advances in neural information processing systems (NeurIPS)*, 2018, pp. 7354–7364.
- [16] F. Ringeval, B. Schuller, M. Valstar, J. Gratch, R. Cowie, S. Scherer, S. Mozgai, N. Cummins, M. Schmitt, and M. Pantic, “Avec 2017: Real-life depression, and affect recognition workshop and challenge,” in *Proc. ACM Multimedia Workshop*, 2017, pp. 3–9.
- [17] S. Chen, Q. Jin, J. Zhao, and S. Wang, “Multimodal multi-task learning for dimensional and continuous emotion recognition,” in *Proc. ACM Multimedia Workshops*, 2017, pp. 19–26.
- [18] T. Dang, B. Stasak, Z. Huang, S. Jayawardena, M. Atcheson, M. Hayat, P. Le, V. Sethu, R. Goecke, and J. Epps, “Investigating word affect features and fusion of probabilistic predictions incorporating uncertainty in avec 2017,” in *Proc. ACM Multimedia Workshops*, 2017, pp. 27–35.
- [19] J. Huang, Y. Li, J. Tao, Z. Lian, Z. Wen, M. Yang, and J. Yi, “Continuous multimodal emotion prediction based on long short term memory recurrent neural network,” in *Proc. ACM Multimedia Workshops*, 2017, pp. 11–18.
- [20] B. Schuller, S. Steidl, A. Batliner, P. Marschik, H. Baumeister, F. Dong, S. Hantke, F. Pokorny, et al., “The interspeech 2018 computational paralinguistics challenge: Atypical & self-assessed affect, crying & heart beats,” 2018, pp. 122–126.
- [21] P. Tzirakis, J. Chen, S. Zafeiriou, and B. Schuller, “End-to-end multimodal affect recognition in real-world environments,” *Information Fusion*, vol. 68, pp. 46–53, 2021.
- [22] G. Trigeorgis, F. Ringeval, R. Brueckner, E. Marchi, Mihalis A. Nicolaou, B. Schuller, and S. Zafeiriou, “Adieu features? end-to-end speech emotion recognition using a deep convolutional recurrent network,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 5200–5204.
- [23] P. Tzirakis, J. Zhang, and B. Schuller, “End-to-end speech emotion recognition using deep neural networks,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5089–5093.
- [24] L. Tarantino, P. Garner, and A. Lazaridis, “Self-attention for speech emotion recognition,” *Proc. Interspeech 2019*, pp. 2578–2582, 2019.
- [25] M. Neumann and N. T. Vu, “Attentive convolutional neural network based speech emotion recognition: A study on the impact of input features, signal length, and acted speech,” *Proc. Interspeech 2017*, pp. 1263–1267, 2017.
- [26] J. Han, Z. Zhang, Z. Ren, and B. Schuller, “Implicit fusion by joint audiovisual training for emotion recognition in mono modality,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 5861–5865.
- [27] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [28] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: a simple way to prevent neural networks from overfitting,” *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.