

Advancing Multiple Instance Learning with Attention Modeling for Categorical Speech Emotion Recognition

Shuiyang Mao¹, P. C. Ching¹, C.-C. Jay Kuo² and Tan Lee¹

¹Department of Electronic Engineering, The Chinese University of Hong Kong, Hong Kong

²Ming Heish Department of Electrical Engineering, University of Southern California, USA

maoshuiyang@link.cuhk.edu.hk, {pcching, tanlee}@ee.cuhk.edu.hk, cckuo@sipi.usc.edu

Abstract

Categorical speech emotion recognition is typically performed as a sequence-to-label problem, i. e., to determine the discrete emotion label of the input utterance as a whole. One of the main challenges in practice is that most of the existing emotion corpora do not give ground truth labels for each segment; instead, we only have labels for whole utterances. To extract segment-level emotional information from such weakly labeled emotion corpora, we propose using multiple instance learning (MIL) to learn segment embeddings in a weakly supervised manner. Also, for a sufficiently long utterance, not all of the segments contain relevant emotional information. In this regard, three attention-based neural network models are then applied to the learned segment embeddings to attend the most salient part of a speech utterance. Experiments on the CASIA corpus and the IEMOCAP database show better or highly competitive results than other state-of-the-art approaches.

Index Terms: categorical speech emotion recognition, weak labeling, multiple instance learning, attention modeling

1. Introduction

Automatic speech emotion recognition (ASER) aims to decode emotional content from audio signals. It has constituted an active research topic in the field of human-machine interaction (HCI). Detection of lies, monitoring of call centers, and medical diagnoses are also considered as promising application scenarios of speech emotion recognition.

Categorical speech emotion recognition at utterance level can be formulated as a sequence-to-label problem. The input utterance is divided into a sequence of acoustic segments, and the output is a single label of emotion type. A few previous studies explored the use of segment units for categorical speech emotion recognition and demonstrated that combining segment-level prediction results led to superior performance [1, 2, 3]. These prior works were mostly based on conventional models, such as *support vector machine* (SVM) and *K-nearest neighbors* (K-NN), for segment classification. In this paper, a *convolutional neural network* (CNN) is applied to extract emotionally relevant features, i. e., to learn emotion relevant segment embeddings. Also, most of the existing emotion corpora do not provide ground truth labels at segment level; instead, they only have labels for whole utterances. A viable solution is to learn local concepts from global annotations, which is the main idea of *multiple instance learning* (MIL) [4, 5]. MIL has been successfully applied to sound event detection (SED) [6, 7], speech recognition [8], and image analysis [5]. In the MIL problem statement, the training set contains labeled bags that comprise many unlabeled instances, and the task is to predict the labels of unseen bags and instances. For categorical speech emotion

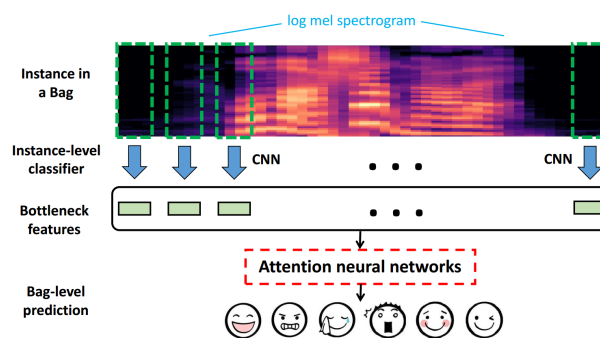


Figure 1: Illustration of the proposed framework for categorical speech emotion recognition.

recognition, each utterance is treated as a bag, and segments within the utterance as instances. One main feature of this work is the application of deep learning of feature representation in the MIL framework to learn segment embeddings from weak labels at utterance level. Compared to the raw features such as MFCC, energy, or pitch, the learned segment embeddings are more tied to the task of interest, thus naturally highlighting salient portions of the data, which we conjecture would offer an advantage in the final classification.

A key question is how to enable a deep learning model to identify and focus on the most salient parts of a speech utterance when making an utterance-level decision with the learned segment embeddings. In this regard, attention neural network models are investigated. The key idea behind the attention mechanism is to align the input-output sequences such that, in the decoding phase, the major contribution is from the corresponding encoded information. In contrast, the effect of irrelevant ones is minimized. In this work, attention modeling is expected to facilitate a structurally meaningful composition of the utterance representation from the learned emotionally relevant segment embeddings. This is the first attempt to combine MIL-based deep learning of segment embedding with attention modeling for categorical speech emotion recognition.

2. Methodology

Figure 1 illustrates a schematic approach of the proposed method. It comprises a CNN model trained to learn emotionally salient segment embeddings from log-Mel filterbank features of individual segments. The learned segment embeddings are then used as inputs for utterance-level emotion recognition, which is achieved by a dense-layer neural network implemented with various attention mechanisms.




Figure 2: An example of CNN outputs for the audio file “Happy_liuchanhg_382.wav” from the CASIA corpus.

2.1. MIL-based segment embedding learning

We formulate our segment-based approach as a MIL problem following the instance space paradigm [4]. Each utterance (bag) is first divided into a sequence of segments (instances). These individual segments are then used to train a CNN model. The learned CNN aims to generate emotionally salient embeddings for each segment.

2.1.1. Log Mel spectrograms

For the segment-level features, we use the 64-bin log Mel filterbanks, which have been extensively evaluated in the existing literature [9, 10, 11]. They are computed by *short-time Fourier transform* (STFT) with a window length of 25 ms, hop length of 10 ms, and FFT length of 512. Subsequently, 64-bin log Mel filterbanks are derived from each short-time frame, and the frame-level features are combined to form a time-frequency matrix representation of the segment.

2.1.2. Segment embeddings

Our segment-based method must address how to train a segment-level model without access to a training set of labeled segments. To address this problem, we follow the most straightforward approach, called Single Instance Learning (SIL) [12], i.e., each segment inherits the label of the utterance where it lies. A CNN is then trained on the resulting dataset. The outputs of the penultimate layer of the trained CNN, which we refer to as *segment embeddings* in this work, are stored and will be employed as inputs to the subsequent recognition part. Besides, a softmax layer sits on top of the CNN model and aims to predict a probability distribution \mathbf{P} as follows:

$$\mathbf{P} = [p(e_1), p(e_2), \dots, p(e_K)]^T \quad (1)$$

where K denotes the number of possible emotions.

Figure 2 shows an example of a probability distribution predicted by the trained CNN for the audio file “Happy_liuchanhg_382.wav” from the CASIA corpus. It can be observed that: (1) the probability distribution of each segment changes across the whole utterance; (2) most of the segments convey information that conforms to the utterance where they lie; and (3) there are segments within one utterance that do not convey any information about the target emotion class or that are more related to other classes, which constitutes confusing information. If we can place additional focus on these more relevant segments, system performance might be improved. In this regard, we have developed three attention-based neural networks, which are described in detail in the following section.

2.2. Attention-based neural networks

Attention neural networks assume that the bag-level prediction can be constructed as a weighted sum of the instance-level predictions. Herein, three attention-based neural network models are investigated and compared, i.e., decision-level single attention (D-Single-Att.) [13], decision-level multiple attention (D-Multi-Att.) [14] and feature-level attention (Feature-Att.) [15], as shown in Figure 3(a)-(c), respectively. We denote the input segment embeddings within a certain speech utterance as $\mathbf{X} \in \mathbb{R}^{T \times M}$, where T is the number of segments and M represents the dimension of segment embeddings. The output of the second fully-connected (FC) layer is denoted as \mathbf{h} , which has a dimension of 120, i.e., H is set to 120 in Figure 3.

2.2.1. Decision-level single attention

In the decision-level single attention model (as shown in Figure 3(a)), an attention function is applied on the predictions of the instances to obtain the bag-level prediction:

$$F(\mathbf{B})_k = \sum_{\mathbf{h} \in \mathbf{B}} w(\mathbf{h})_k f(\mathbf{h})_k \quad (2)$$

where k denotes the k -th emotion class of the instance-level prediction $f(\mathbf{h}) \in [0, 1]^K$ and the bag-level prediction $F(\mathbf{B}) \in [0, 1]^K$, and $w(\mathbf{h})_k \in [0, 1]$ is a weight of $f(\mathbf{h})_k$ that we refer to as a *decision-level attention function*:

$$w(\mathbf{h})_k = s(\mathbf{h})_k / \sum_{\mathbf{h} \in \mathbf{B}} s(\mathbf{h})_k \quad (3)$$

where $s(\cdot)$ can be any non-negative function (i.e., Softmax non-linearity) to ensure that attention $w(\cdot)$ is normalized. Both the attention function $w(\cdot)$ and the instance-level classifier $f(\cdot)$ depend on a set of learnable parameters.

2.2.2. Decision-level multiple attention

The decision-level multiple attention model is an extension of the above decision-level single attention model. It consists of several single attention modules (we herein use two attention modules, as shown in Figure 3(b)) applied to intermediate neural network layers. The outputs of these attention modules are concatenated.

2.2.3. Feature-level attention

The limitation of the above decision-level attention neural networks is that the attention function $w(\cdot)$ is only applied to the prediction of the instances $f(\mathbf{h})$. To address this constraint, we also investigate the effect of applying the attention function to the outputs of the hidden layers, which we refer to as the feature-level attention (as shown in Figure 3(c)), in which the bag-level representation \mathbf{U} can be modeled as:

$$\mathbf{U}_d = \sum_{\mathbf{h} \in \mathbf{B}} v(\mathbf{h})_d q(\mathbf{h})_d \quad (4)$$

where d denotes the d -th dimension of the hidden layer output $q(\mathbf{h}) \in \mathbb{R}^D$ and the bag-level representation $\mathbf{U} \in \mathbb{R}^D$; and $v(\mathbf{h})_d \in [0, 1]$ is a weight of $q(\mathbf{h})_d$ that we refer to as a *feature-level attention function*:

$$v(\mathbf{h})_d = u(\mathbf{h})_d / \sum_{\mathbf{h} \in \mathbf{B}} u(\mathbf{h})_d \quad (5)$$

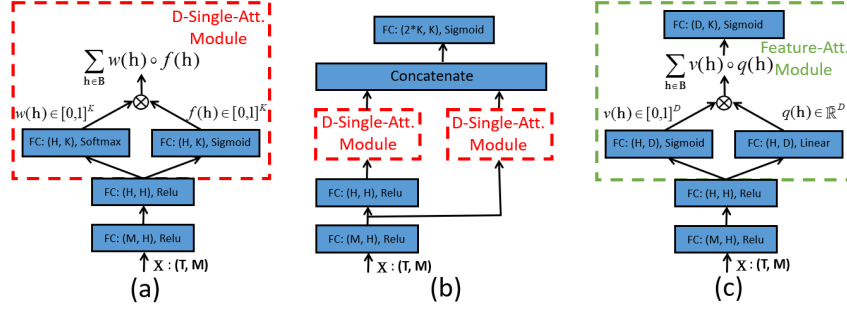


Figure 3: (a) Decision-level single attention neural network; (b) decision-level multiple attention neural network; (c) feature-level attention neural network. (\odot : Hadamard product; \sum : element-wise summation; T : length of input sequence; M : dimension of input bottleneck features; H : dimension of FC layer; K : number of emotion classes; D : dimension of feature-level attention function.)

where $u(\cdot)$ can be any non-negative function (i.e., Sigmoid non-linearity) to ensure that attention $v(\cdot)$ is normalized. Both the attention function $v(\cdot)$ and the instance-level feature mapping function $q(\cdot)$ depend on a set of learnable parameters. The prediction of a bag \mathbf{B} can then be obtained by classifying the bag-level representation \mathbf{U} as follows:

$$F(\mathbf{B}) = g(\mathbf{U}) \quad (6)$$

where $g(\cdot)$ is the final classifier that corresponds to the last neural network layer.

3. Emotion Corpora

Two different emotion corpora are used to evaluate the validity of the proposed method, namely, a Chinese emotional corpus (CASIA) [16] and an English emotional database (IEMOCAP) [17], which have been extensively evaluated in the literature.

Specifically, the CASIA corpus [16] contains 9,600 utterances that are simulated by four subjects (two males and two females) in six different emotional states, i.e., angry, fear, happy, neutral, sad, and surprise. In our experiments, we only use 7,200 utterances that correspond to 300 linguistically neutral sentences with the same statements. All of the emotion categories are selected.

The IEMOCAP database [17] was collected using motion capture and audio/video recording over five dyadic sessions with 10 subjects. At least three evaluators annotated each utterance in the database with the categorical emotion labels chosen from the set: angry, disgusting, excited, fear, frustrate, happy, neutral, sad, surprise, and others. We consider only the utterances with majority agreement (i.e., at least two out of three evaluators assigned the same emotion label) over the emotion classes of angry, happy (combined with the “excited” category), neutral and sad, which results in 5,531 utterances in total.

4. Experiments

4.1. Setup

In our experiment, the size of each speech segment is set to 32 frames, i.e., the total length of a segment is $10 \text{ ms} \times 32 + (25 - 10) \text{ ms} = 335 \text{ ms}$, shifting 60 ms each time. In this way, we collected approximately 200,000 segments for the CASIA corpus and 300,000 segments for the IEMOCAP database, respectively. Moreover, since the input length for our attention neural networks has to be equal for all samples, we heuristically set the maximal length for each speech utterance to the average

duration of each dataset, i.e., 2.07 s for CASIA and 4.55 s for IEMOCAP, respectively. Longer speech utterances are cut at the maximal length, and shorter ones are padded with zeros.

The architecture of the CNN model is similar to the *SegCNN* model as used in our previous work [18]. The only change we made was to the last three FC layers, i.e., $\{256, 64, K\}$ units, respectively, where 64 and K correspond to the dimension of segment embeddings and the number of possible emotions, respectively. In the training stage, for both CNN and attention neural networks, ADAM [19] optimizer with the default setting in Tensorflow [20] was used, with an initial learning rate of 0.001 and an exponential decay scheme with a rate of 0.8 every two epochs. The batch size was set to 128. Early stopping with patience of 3 epochs was utilized to mitigate an overfitting problem.

For the CASIA corpus, we perform leave-one-fold-out ten-fold cross-validation experiments. For the IEMOCAP database, the leave-one-session-out five-fold cross-validation method is carried out. For both datasets, a second cross-validation is performed since we need to utilize the segment-level results to train the attention neural networks. The results are presented in terms of unweighted accuracy (UA).

4.2. Baseline systems

4.2.1. Maxout segment embedding (CNN-MAX-RF)

In the CNN-MAX-RF baseline, the maxout pooling is directly applied on the segment embedding $\mathbf{X} \in \mathbb{R}^{T \times M}$:

$$\mathbf{U}_m = \max_{0 \leq t \leq T} \{\mathbf{X}_m\} \quad (7)$$

A *Random Forest* (RF) is then used to make the utterance-level prediction based on the resultant utterance representation \mathbf{U} .

4.2.2. Averaging segment embedding (CNN-AVG-RF)

The CNN-AVG-RF baseline is similar to the above CNN-MAX-RF baseline. The only difference is that the maxout pooling in the CNN-MAX-RF baseline is replaced by an average pooling in the CNN-AVG-RF baseline.

4.2.3. Maxout pooling (CNN-MP)

In the CNN-MP baseline, the maxout pooling is applied on the instance-level prediction $f(\mathbf{h}) \in [0, 1]^K$ across a certain speech utterance to obtain the bag-level prediction:

$$F(\mathbf{B})_k = \max_{\mathbf{h} \in \mathbf{B}} \{f(\mathbf{h})_k\} \quad (8)$$

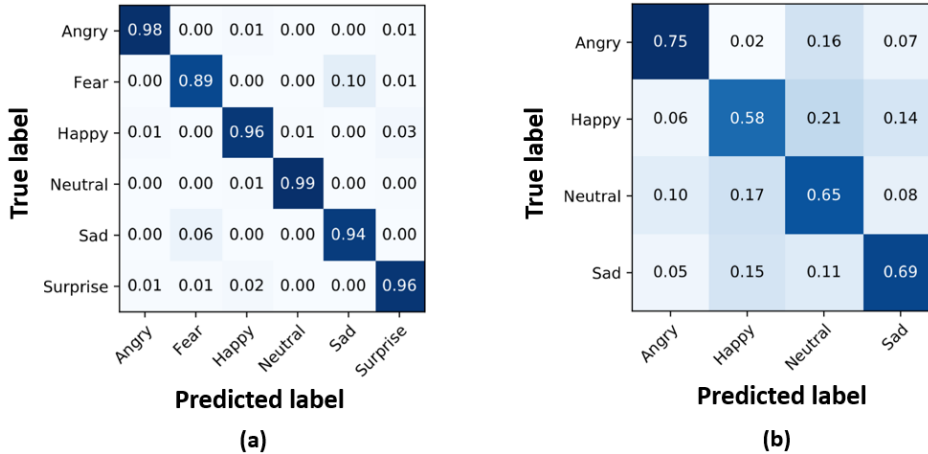


Figure 4: Confusion matrices obtained using feature-level attention modeling for: (a) CASIA corpus and (b) IEMOCAP database.

Table 1: Comparison of UAs on the CASIA corpus.

Methods for comparison	UA [%]
ELM-Decision Tree [21]	89.60
DNN-HMM [22]	91.32
LSTM-TF-Att. [23]	92.80
CNN-MAX-RF (<i>baseline</i>)	91.10
CNN-AVG-RF (<i>baseline</i>)	91.75
CNN-MP (<i>baseline</i>)	92.11
CNN-AP (<i>baseline</i>)	92.36
CNN-D-Single-Att. (<i>ours</i>)	93.75
CNN-D-Multi-Att. (<i>ours</i>)	94.57
CNN-Feature-Att. (<i>ours</i>)	95.32

Table 2: Comparison of UAs on the IEMOCAP database.

Methods for Comparison	UA [%]
CNN-LSTM [10]	59.40
DNN-HMM [22]	62.23
FCN-Att. [24]	63.90
CNN-MAX-RF (<i>baseline</i>)	61.08
CNN-AVG-RF (<i>baseline</i>)	61.40
CNN-MP (<i>baseline</i>)	62.11
CNN-AP (<i>baseline</i>)	62.74
CNN-D-Single-Att. (<i>ours</i>)	63.34
CNN-D-Multi-Att. (<i>ours</i>)	65.82
CNN-Feature-Att. (<i>ours</i>)	66.74

4.2.4. Average pooling (CNN-AP)

Similarly, for the CNN-AP baseline, the average pooling is applied on the instance-level prediction $f(\mathbf{h}) \in [0, 1]^K$ across a particular speech utterance to obtain the bag-level prediction:

$$F(\mathbf{B})_k = \text{mean}_{\mathbf{h} \in \mathbf{B}} \{f(\mathbf{h})_k\} \quad (9)$$

4.3. Results and analysis

Tables 1-2 show the experimental results on the two mentioned emotion corpora, respectively. The following can be seen: (1) our baseline systems achieved respectable results on both datasets, which proved the effectiveness of the MIL-based framework; (2) the last two baselines (i. e., the CNN-MP baseline and the CNN-AP baseline) consistently outperformed the first two baselines (i. e., the CNN-MAX-RF baseline and the CNN-AVG-RF baseline). This performance gain might derive from the joint optimization of the aggregation strategy of segment-level representations and the utterance-level decision making of the last two baselines; (3) the attention-based methods substantially augmented the performance of the baseline systems overall. This is mainly attributed to the effectiveness of the attention modeling; (4) due to the positive combination of different attention modules, the decision-level multiple attention modeling achieved noticeably better performance than the decision-level single attention modeling on both datasets; (5)

the feature-level attention modeling outperformed the decision-level attention neural networks by a significant margin. This is due to the fact that the dimension of $v(\mathbf{h})$ (i. e., D) can be any value, while the dimension of $w(\mathbf{h})$ is fixed to be the number of emotion classes K . With an increase in the dimension of $v(\mathbf{h})$, the capacity of feature-level attention neural networks is increased; and (6) for the CASIA corpus, our feature-level attention-based system achieved the highest recognition accuracy of 95.32%, establishing a new benchmark (to the best of our knowledge). For the IEMOCAP database, which might constitute a more challenging dataset, our methods also achieved competitive results. Figure 4 shows the confusion matrices obtained using feature-level attention modeling on both datasets, respectively.

5. Conclusion

In this paper, we proposed to combine multiple instance learning with attention neural networks for better modeling of categorical speech emotion recognition. Three attention-based neural network models were investigated and compared. Experimental results on two well-known emotion corpora showed competitive outcomes. Since we herein blindly used all segments to train the segment-level classifier, it is anticipated with proper segment selection strategy, better results are expected. More advanced neural network architectures and better algorithm optimization will also be investigated in the near future.

6. References

- [1] J. H. Jeon, R. Xia, and Y. Liu, "Sentence level emotion recognition based on decisions from subsentence segments," in *Proc. ICASSP*, 2011, pp. 4940–4943.
- [2] M. T. Shami and M. S. Kamel, "Segment-based approach to the recognition of emotions in speech," in *Proc. ICME*, 2005, p. 4.
- [3] B. Schuller and G. Rigoll, "Timing levels in segment-based speech emotion recognition," in *Proc. INTERSPEECH*, 2006, pp. 1818–1821.
- [4] J. Amores, "Multiple instance classification: Review, taxonomy and comparative study," *Artificial Intelligence*, vol. 201, pp. 81–105, 2013.
- [5] Y. Xu, T. Mo, Q. Feng, P. Zhong, M. Lai, I. Eric, and C. Chang, "Deep learning of feature representation with multiple instance learning for medical image analysis," in *Proc. ICASSP*, 2014, pp. 1626–1630.
- [6] T.-W. Su, J.-Y. Liu, and Y.-H. Yang, "Weakly-supervised audio event detection using event-specific gaussian filters and fully convolutional networks," in *Proc. ICASSP*, 2017, pp. 791–795.
- [7] A. Kumar and B. Raj, "Audio event detection using weakly labeled data," in *Proc. ACM International Conference on Multimedia*, 2016, pp. 1038–1047.
- [8] Y. Wang, J. Li, and F. Metze, "Comparing the max and noisy-or pooling functions in multiple instance learning for weakly supervised sequence learning tasks," in *Proc. INTERSPEECH*, 2018, pp. 1339–1343.
- [9] Z. Huang, M. Dong, Q. Mao, and Y. Zhan, "Speech emotion recognition using cnn," in *Proc. ACM International Conference on Multimedia*, 2014, pp. 801–804.
- [10] A. Satt, S. Rozenberg, and R. Hoory, "Efficient emotion recognition from speech using deep learning on spectrograms," in *Proc. INTERSPEECH*, 2017, pp. 1089–1093.
- [11] L. Zhang, L. Wang, J. Dang, L. Guo, and H. Guan, "Convolutional neural network with spectrogram and perceptual features for speech emotion recognition," in *Proc. ICONIP*, 2018, pp. 62–71.
- [12] R. C. Bunescu and R. J. Mooney, "Multiple instance learning for sparse positive bags," in *Proc. ICML*, 2007, pp. 105–112.
- [13] Q. Kong, Y. Xu, W. Wang, and M. D. Plumbley, "Audio set classification with attention model: A probabilistic perspective," in *Proc. ICASSP*, 2018, pp. 316–320.
- [14] C. Yu, K. S. Barsim, Q. Kong, and B. Yang, "Multi-level attention model for weakly supervised audio classification," in *arXiv:1803.02353*, 2018.
- [15] Q. Kong, C. Yu, Y. Xu, T. Iqbal, W. Wang, and M. D. Plumbley, "Weakly labelled audioset tagging with attention neural networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 11, pp. 1791–1802, 2019.
- [16] J. Tao, F. Liu, M. Zhang, and H. Jia, "Design of speech corpus for mandarin text to speech," in *Proc. the 4th Workshop on Blizzard Challenge*, 2005.
- [17] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "Iemocap: Interactive emotional dyadic motion capture database," *Language resources and evaluation*, vol. 42, no. 4, p. 335, 2008.
- [18] S. Mao, P. C. Ching, and T. Lee, "Deep learning of segment-level feature representation with multiple instance learning for utterance-level speech emotion recognition," in *Proc. INTERSPEECH*, 2019, pp. 1686–1690.
- [19] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [20] M. Abadi *et al.*, "Tensorflow: A system for large-scale machine learning," in *Proc. OSDI*, 2016, pp. 265–283.
- [21] Z.-T. Liu, M. Wu, W.-H. Cao, J.-W. Mao, J.-P. Xu, and G.-Z. Tan, "Speech emotion recognition based on feature selection and extreme learning machine decision tree," *Neurocomputing*, vol. 273, pp. 271–280, 2018.
- [22] S. Mao, D. Tao, G. Zhang, P. C. Ching, and T. Lee, "Revisiting hidden markov models for speech emotion recognition," in *Proc. ICASSP*, 2019, pp. 6715–6719.
- [23] Y. Xie, R. Liang, Z. Liang, C. Huang, C. Zou, and B. Schuller, "Speech emotion classification using attention-based lstm," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 11, pp. 1675–1685, 2019.
- [24] Y. Zhang, J. Du, Z. Wang, J. Zhang, and Y. Tu, "Attention based fully convolutional network for speech emotion recognition," in *Proc. APSIPA ASC*, 2018, pp. 1771–1775.