



SpeechEQ: Speech Emotion Recognition based on Multi-scale Unified Datasets and Multitask Learning

Zuheng Kang, Junqing Peng, Jianzong Wang*, Jing Xiao

Ping An Technology (Shenzhen) Co., Ltd.

{kangzuheng896, pengjq, wangjianzong347, xiaojing661}@pingan.com.cn

Abstract

Speech emotion recognition (SER) has many challenges, but one of the main challenges is that each framework does not have a unified standard. In this paper, we propose SpeechEQ, a framework for unifying SER tasks based on a multi-scale unified metric. This metric can be trained by Multitask Learning (MTL), which includes two emotion recognition tasks of Emotion States Category (EIS) and Emotion Intensity Scale (EIS), and two auxiliary tasks of phoneme recognition and gender recognition. For this framework, we build a Mandarin SER dataset - SpeechEQ Dataset (SEQD). We conducted experiments on the public CASIA and ESD datasets in Mandarin, which exhibit that our method outperforms baseline methods by a relatively large margin, yielding 8.0% and 6.5% improvement in accuracy respectively. Additional experiments on IEMOCAP with four emotion categories (i.e., angry, happy, sad, and neutral) also show the proposed method achieves a state-of-the-art of both weighted accuracy (WA) of 78.16% and unweighted accuracy (UA) of 77.47%.

Index Terms: Speech Emotion Recognition, Multi-task Learning, Phoneme Recognition, Gender Recognition

1. Introduction

Emotions, an “implicit channel” that transmits explicit messages [1], play a vital role in human communication. Speech is an important carrier of emotion and the easiest way to record information completely. Modern people usually communicate by phone or through APPs. Unfortunately, existing technologies lack standards for speech emotion information. If designed a unified SER framework that automatically recognizes the human emotional state categories (ESC), such as happy, sad, anger, and its emotional intensity scales (EIS) expressed in natural speech, the intelligence systems can better perceive human thoughts and have better interactivity.

Recently, deep-learning-based approaches show great performance in extracting hidden speech information, including age [2], facial expression [3] and emotion, etc.

The improvement of network architecture with recurrent neural networks (RNN) and attention mechanism [4, 5], or analyzing and fusing the textual modal information [6, 7] has demonstrated impressive results in solving SER problem. However, since emotion is a complex psychological process, a single task may not be sufficient to capture complicated emotional features.

Using a transfer model for different tasks can also improve accuracy. Because it essentially uses external knowledge to make the model more robust. To capture emotional information, researchers combine the power of automatic speech recognition (ASR) with natural language processing (NLP) into the SER

network structure to improve model performance [8, 9]. However, their solutions are to fuse the information of different tasks in the last few layers (high-level features) of the network, which makes the information of different tasks difficult to communicate with each other. In fact, the information of different tasks is intrinsically correlated with each other. In SER, the speed and the way of speaking will affect the expression of human emotions. Therefore, a unified model will comprehensively consider these factors to give a reasonable inference.

Multi-task learning (MTL) uses a shared backbone model to simultaneously optimize multiple objectives in different tasks. The advantage comes from side information and cross-regularization of different tasks. At the same time, joint optimization also brings challenges [10]. In SER, MTL provides an idea for the model to learn multimodal information simultaneously [9, 11, 12]. However, these methods use only one dataset for training and testing, which leads to a lack of model generalization.

Unsupervised learning utilizes multiple datasets to generate high-quality speech features [13, 14] for SER. However, this method does not fully exploit hidden information in these datasets.

The released SER datasets have: (1) content-dependent (such as IEMOCAP [15]) or content independent (such as RAVDESS [16], ESD [17], CREMA-D [18]); (2) contain EIS (such as IEMOCAP, CREMA-D, RAVDESS) or not (CASIA [19]); (3) lack a unified standard for the ESC in each dataset.

To find this unifying criterion, we took inspiration from psychologists’ theories of emotion. One of the most well-known theories is the emotional wheel proposed by Psychologist Robert Plutchik. He theorized 8 emotions with 24 “primary”, “secondary”, and “tertiary” dyads [20]. Later, E. Cambria et al. [21] updated his theory and created the hourglass of emotions - he divided emotion into 4 dimensions, but reversed the position of each emotion, such that the intensity of emotion on each dimension could be a numerical value. These theories give us an idea to create a calculable metric for speech emotions. And the main contributions of our works are as follows:

- This paper proposed SpeechEQ, using a multi-scale unified metric, SpeechEQ Metric (SEQM), that unifies all frameworks. It can be trained with an MTL framework to simultaneously perform two emotion recognition tasks, ESC and EIS, and two auxiliary tasks of phoneme recognition and gender recognition;
- We build a Mandarin SER dataset - SpeechEQ dataset (SEQD) to demonstrate that using this metric, the emotion recognition accuracy of each Mandarin SER dataset can be improved;
- The effectiveness of this method is also verified on the public English SER dataset IEMOCAP;

*Corresponding author: Jianzong Wang, jzwang@188.com

2. Methodology

2.1. Multiscale Unified Metric

In order to unify different frameworks into a unified standard, we need a metric. Inspired by the emotional hourglass theory, we divide human emotions into 8 ESC. Then, use a numerical value to represent the intensity of the emotion: a value from 1 to 4 represents the EIS from low to high. Finally, very weak emotions with low EIS values ranging from 0 to 1 are defined as the 9th category – Neutral. We name it, the SpeechEQ metric (SEQM), shown in Figure 1. With this metric, multiple datasets can be fused into a Multiscale Unified Dataset (MsUD) – a label unification mechanism.

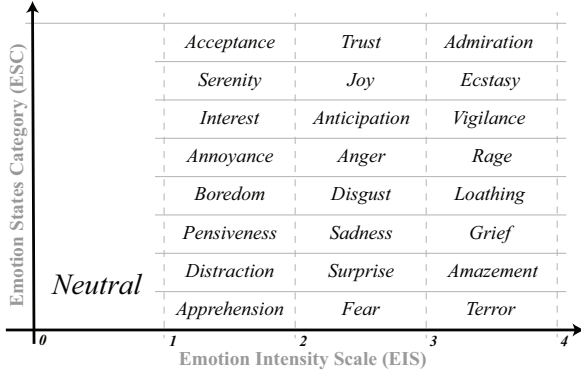


Figure 1: the SpeechEQ Metric for SER.

In modeling, since emotion categories are independent of each other, but the discrimination of emotional intensity is vague, sometimes indistinguishable, our proposed model will classify emotions into 9 categories and regress emotional intensity. To create the gold label for training, if the EIS is divided into 3 levels (i.e., low, medium, high), the gold label will be set to 1.5, 2.5, and 3.5 (neutral is 0). If with 2 levels, set them to 2 and 3 respectively. If emotional intensity has been recorded in other forms, rescale to the range of 1 to 4.

2.2. SpeechEQ Dataset

To test the effectiveness of the SEQM, we constructed the SpeechEQ Dataset (SEQD), a content-independent Mandarin SER dataset, to help us improve the overall model performance in Mandarin SER tasks. This dataset contains a total of 2.3 hours of speech, 1648 audio clips from 20 speakers (10 males and 10 females) at a sampling rate of 16kHz and a precision of 16 bits. It was recorded with a Huawei phone in a medium-sized conference room at a signal-to-noise ratio (SNR) of about 20dB. The speaker is about 20cm from the microphone, sitting in the center of the room. See Table 1 for details (Neutral is 73).

Table 1: Number of utterances in SEQD with emotions and level (use its Medium emotion name to represent overall ESC).

	Trust	Joy	Anticipation	Anger	Disgust	Sadness	Surprise	Fear
Low	62	59	60	72	66	61	61	72
Medium	62	63	65	77	73	77	63	63
High	70	59	60	71	67	62	70	60

To build this dataset, each speaker independently wrote 3 to 5 sentences describing the emotion for each of the 25 emotions in the SEQM and performed it independently in the tone

of talking to someone in everyday conversation. Then, each utterance was judged independently by three judges. If more than one judge judges the recorded speech emotion to be inaccurate, the speaker is notified to re-record until all three judges agree that the emotion is matched.

2.3. Model Structure

For the speech feature extraction, the structure of ECAPA-TDNN [22] has efficient design structures such as Res2Net [23] and Squeeze Excitation blocks (SE) [24]. However, it is originally for the speaker recognition task, which results in a smaller receptive field. SER requires more frames to understand contextual information for more comprehensive reasoning, which needs a larger receptive field.

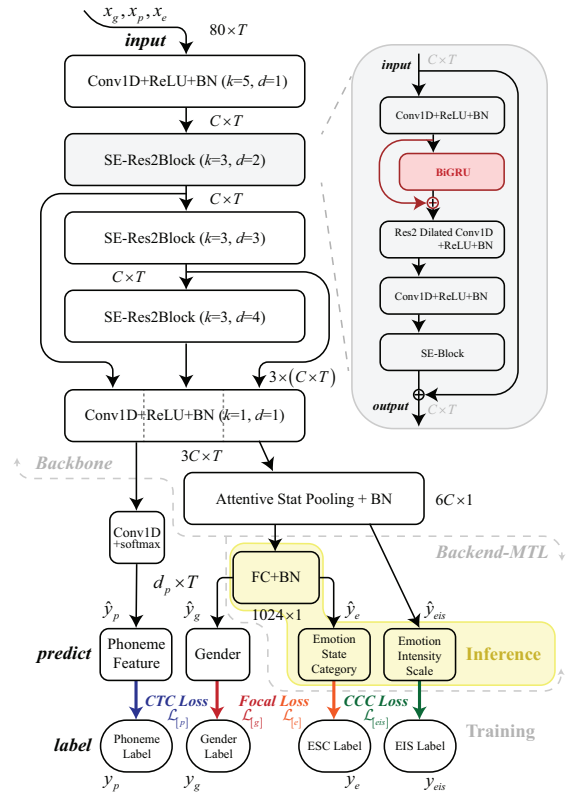


Figure 2: Network topology of the SpeechEQ framework.

In Figure 2, in the backbone model, all structures are the same with ECAPA except Res-BiGRU (the residual structure of bidirectional GRU) in the red background, which has been added to SE-Res2Block that extends the receptive field. The backend model is changed to the structure of MTL. In our SER framework, in addition to predicting the category and intensity of emotions, the model also predicts the phonemes in speech and the gender of the speaker as auxiliary tasks.

2.4. Training and Inference

In the training phase, the generated predictions can have access to the labels via blue, red, orange, and green paths to generate loss functions respectively. For each utterance in the training set, the blue path has access to the gold phoneme sequences labels y_p translated from text labels, the red path to the ground-truth gender labels y_g , the orange path to the gold ESC labels

y_e and the green path to the gold EIS labels y_{eis} .

For phoneme sequence prediction (blue task), the Connectionist Temporal Classification (CTC) [25] as the metric between predicted phoneme probability vectors and the given gold phoneme sequence labels. Assuming that the number of frames of speech features is T , the phoneme feature dimension is d_p . Thus, the CTC loss for phoneme recognition task is:

$$\mathcal{L}_{[p]} = \text{CTC}(y'_p, y_p), \text{ where } y'_p = \text{softmax}(\hat{y}_p) \in \mathcal{P}^{T \times d_p} \quad (1)$$

For classification of emotion and gender (red and orange task), the proposed classification tasks are trained with a focal loss (FL) [26] with a tunable parameter γ . Assume the gender and ESC feature dimension is d_g (=2) and d_e (=9). The loss for ESC and gender recognition tasks is in Equation 2, the sign * denotes either g for gender, or e for ESC:

$$\mathcal{L}_{[*]} = \text{FL}(\hat{y}_*, y_*, \gamma), \text{ where } \hat{y}_* \in \mathcal{P}^{d_*} \quad (2)$$

For regression of EIS (green task), the Lin’s Concordance Correlation Coefficient (CCC) [27] is used. Let $\mathcal{L}_{[eis]}$ denote the loss function for EIS:

$$\mathcal{L}_{[eis]} = 1 - \text{CCC}(\hat{y}_{eis}, y_{eis}), \text{ where } \hat{y}_{eis} \in \mathcal{P} \quad (3)$$

We introduce a hyper-parameter α, β, η to combine four losses into a single one \mathcal{L} . α controls the relative importance of the CCC loss for EIS, β of CTC loss for phoneme recognition, η of classification loss for gender recognition. Finally, the model will be optimized with the following objective w.r.t θ :

$$\min_{\theta} \mathcal{L} = \mathcal{L}_{[e]} + \alpha \mathcal{L}_{[eis]} + \beta \mathcal{L}_{[p]} + \eta \mathcal{L}_{[g]} \quad (4)$$

In the training process, the ASR dataset, gender dataset, and emotion dataset are respectively converted to speech features and sent to the network, and their corresponding losses are calculated separately. Finally, combine these losses into one loss to optimize the entire model.

However, in EIS tasks, some data do not have labels assigned. For this type of data, here’s a label ignore mechanism for unspecified EIS labels by setting a mask, i.e., temporarily setting them to a fake value of -1 when building the gold labels. During training, these values can be captured and masked by replacing any label value in the mini-batch satisfied $\hat{y}_{eis} < 0$ with the predicted such that there is no error between the predicted value and the gold label value. To prevent the value from exceeding the value range, a clipping function is designed to clip the value to the range of 0 to 4, $y_{eis} = \text{clip}(\hat{y}_{eis})$. The pseudo-code for the detailed training process is in Algorithm 1.

At inference time, two auxiliary tasks are abandoned, and only emotion recognition tasks with ESC and EIS are retained (yellow background in Figure 2). Next, the emotion category softmax can be replaced by the argmax operator, and select the most probable emotion class label as an output.

2.5. Data Augmentation

Since ASR data are recorded in a neutral tone at a moderate speech rate, yet the SER data contains various speech rates and intonations, we will use the package of librosa [28] to adjust the pitch (randomly select pitch parameter from -3 to 3) and the speed (randomly select speed parameter from 0.8 to 1.3) of each utterance during training. Additionally, all data used for training will be augmented by probabilistically adding noise with an SNR ranging from 10dB to 25dB, and reverberation.

3. Experiments

In this paper, a total of 2 parts of the experiment will be performed. The purpose of experiment 1 is to verify whether the

Algorithm 1 SpeechEQ Training Policy

Input: Speech feature inputs x_g, x_p, x_e from gender, phoneme, and emotion datasets respectively
Output: Corresponding predicted gender, phoneme, and ESC, and EIS labels y_g, y_p, y_e, y_{eis}

- 1: initialize model parameters $M(\theta; x)$
- 2: **for** $n \in [1, \text{end}]$ **do**
- 3: **for** i^{th} mini-batch in dataLoader **do**
- 4: prepare input speech features $x_g^{(i)}, x_p^{(i)}, x_e^{(i)}$
- 5: prepare output gold labels $y_g^{(i)}, y_p^{(i)}, y_e^{(i)}, y_{eis}^{(i)}$
- 6: get predictions from multi-datasets:
- 7: $\hat{y}_g^{(i)} \leftarrow M(\theta; x_g^{(i)})$,
- 8: $\hat{y}_p^{(i)} \leftarrow M(\theta; x_p^{(i)})$,
- 9: $\hat{y}_e^{(i)}, \hat{y}_{eis}^{(i)} \leftarrow M(\theta; x_e^{(i)})$
- 10: /* label ignore mechanism */
- 11: **for** j^{th} element in mini-batch **do**
- 12: **if** $\hat{y}_{eis}^{(i,j)} < 0$ **then**
- 13: $y_{eis}^{(i,j)} \leftarrow \text{clip}(\hat{y}_{eis}^{(i,j)})$
- 14: **end if**
- 15: **end for**
- 16: get FL loss for gender $\mathcal{L}_{[g]}$ using Equation 2
- 17: get CTC loss for phoneme $\mathcal{L}_{[p]}$ using Equation 1
- 18: get FL loss for ESC $\mathcal{L}_{[e]}$ using Equation 2
- 19: get CCC loss for EIS $\mathcal{L}_{[eis]}$ using Equation 3
- 20: merge losses with \mathcal{L} using Equation 4
- 21: create and apply gradients with optimizer:
- 22: update model parameters θ ;
- 23: **end for**
- 24: **end for**

Table 2: Dataset details, abbreviation: N-neutral, A-angry, H-happy, S-sad, F-fear, D-disgust, Sr-surprise, Exc-excited.

Dataset	Used Labels	Level	Language
SEQD	SEQM emotions	3 levels	Mandarin
CASIA	A, F, H, N, S, Sr	unspecified	Mandarin
ESD	A, H, N, S	unspecified	Mandarin+English
RAVDESS	A, H, N, S	2 levels	English
CREMA-D	A, H, N, S	3 levels+unspecified	English
IEMOCAP	A, H(+Exc), N, S	continuous	English

MsUD and MTL are helpful to improve the accuracy of each dataset in Mandarin. Experiment 2 confirms that our method is also helpful to the public English dataset – IEMOCAP.

3.1. Datasets and Metrics

Multiple datasets will be used in these experiments, shown in Table 2. The dataset for gender recognition uses the merging of all used SER datasets (training set). For phoneme and emotion recognition tasks, they are determined by different experiments.

Experiment 1: For the CASIA dataset, we perform a leave-one-fold-out ten-fold cross-validation experiment. For the SEQD dataset, we give a fixed test set of 155 utterances. The ESD dataset uses its given test set (Mandarin part, except Surprise). The MsUD is built by the training part of SEQD, CASIA, ESD (Mandarin part, except Surprise). The phoneme recognition training for Mandarin uses aishell-1 [29]. The phoneme gold label is built by the given transcript translated into INITIALS and FINAL_TONE3 by the package pypinyin, then converted to an integer through the phoneme dictionary (dictionary size is 202, with 21 initials, 36 finals, 5 tones, 1 silence). The test metrics for the ESC are weighted accuracy (WA), unweighted accuracy (UA), and the EIS uses the mean square error (MSE). The unweighted accuracy of important emotions (UAI) with neutral, happy, angry, and sad maybe also

Table 3: Ablation study of the proposed model on Mandarin datasets compared with baseline model.

Experiment	CASIA	ESD	SEQD			
	Acc	Acc	WA	UA	UIi	MSE
Baseline	93.27%	89.43%	55.83%	53.87%	64.24%	1.04
w/o MTL w/o MsUD	88.52%	86.76%	46.71%	45.04%	52.64%	0.98
SpeechEQ	96.45%	93.25%	66.14%	65.25%	86.91%	0.68
w/o MTL	95.32%	92.33%	62.47%	61.56%	85.45%	0.67
w/o MsUD	93.15%	90.46%	52.24%	51.30%	67.20%	0.73
w/o MTL w/o MsUD	92.58%	89.67%	51.18%	50.16%	55.01%	0.71

listed. If the number of test samples for each class is the same (WA=UA), we denote it as Acc. Finally, the average score of each cross-validation result is taken as the final performance score. For datasets that do not have EIS labels (i.e., CASIA and ESD), it is trained using the label ignore mechanism.

Experiment 2: For IEMOCAP, we select 4 emotions (i.e., angry, happy+excited, sad, and neutral) – Table 2. The experimental setup is the same as the cited system in Table 4. We use the “dominance” value as its EIS, and rescale them from a range of 1 to 5, to a range of 1 to 4. Then, randomly split the train/test into a 4:1 ratio and perform a 5-fold cross-validation. The MsUD is built by important emotions of RAVDESS, CREMA-D, ESD (English part). The phoneme recognition for English uses Librispeech [30]. Transcripts can be directly translated into phonemes as gold labels via the dictionary CMUDict [31] (dictionary size is 85, with 84 phonemes and 1 silence). Unspecified EIS labels are trained with the label ignore mechanism. The test metrics are the same as in Experiment 1.

3.2. Hyper-parameters

For the speech feature, all datasets will be resampled to the sample rate at 16kHz and precision of 16 bits, and use the 80-dimensional Mel Filter-Banks feature extracted using the librosa package with frame size 25ms, hop size 10ms with Hamming window. The channel parameter C in the convolutional layers for the proposed network is 256. The dimension of the bottleneck in the SE-Block and the attention module is set to 128. The scale dimensions in the Res2Block are set to 8. The tunable parameter γ in FL is 10.

In the first stage, the parameters α , β , η of the combined loss are initially set to be 1, 0.1, 0.1 respectively for fast convergence of each task. All networks are optimized with the Adam optimizer [32] with a learning rate of 1e-4 and weight-decay of 1e-5. The mini-batch size for training is 32.

In the fine-tune stage, both of combined loss parameters β and η gradually decreased from 0.1 to 0.01. Furthermore, the learning rate decreased from 1e-4 to 1e-6.

3.3. Evaluation Results

Experiment 1: The baseline model of CNN+RNN+Attention shares the same backbone model with [11]. Ablation studies are performed on the SpeechEQ framework with or without MTL or MsUD compared to the baseline model.

The results are presented in Table 3, and the EIS value distribution of Neutral, Low, Medium and High for training and testing set of SpeechEQ are shown in the Figure 3a, 3b. Experimental results demonstrate that the feature extractor of SpeechEQ outperforms the baseline model with or without the help of MsUD and MTL. Although adding auxiliary tasks of gender recognition and phoneme recognition improves the accuracy, training models on only a single dataset have limited im-

Figure 3: EIS value distributions of Neutral, Low, Medium and High for training and testing set on SEQD (for SpeechEQ).




Table 4: Ablation study of the proposed model on IEMOCAP compared with cited SER results. The MSE for EIS at SpeechEQ and “w/o MTL w/o MsUD” are 0.74 and 0.81 respectively.

Model	WA	UA
CTC+Attention [33]	67%	69%
BiGRU [34]	66.6%	70.5%
MTL+Attention [12]	76.4%	70.1%
Head Fusion [35]	76.18%	76.36%
SpeechEQ (ours)	78.16%	77.47%
w/o MTL w/o MsUD	73.20%	74.62%

provement. Without the use of MTL, the accuracy of the model is significantly improved after adding a MsUD. Combining the two approaches, the accuracy of the model is further improved on both models. From this result, it can be seen that training the model with either MsUD or MTL, or their combination, improves the accuracy of each dataset. Meanwhile, the regression task of EIS will also be more convergent.

Experiment 2: Compares the evaluation accuracy of training with or without MsUD and MTL, as well as other cited systems trained and tested only with IEMOCAP (results copied directly from the reference) in Table 4. It is proved that the SpeechEQ framework is also applicable to the English dataset: using MsUD (with external data and labels) and MTL can significantly improve the ESC accuracy of the model, and further improve the state-of-the-art results by about 2%. At the same time, the EIS regression task is more convergent.

4. Conclusions

In this paper, we propose SpeechEQ, a framework for unifying emotion recognition tasks: (1) it unifies all SER frameworks with a multi-scale unified metric - SpeechEQ Metric; (2) it can be trained on an MTL framework to perform not only ESC classification and EIS regression, but also auxiliary tasks of gender recognition and phoneme recognition. Subsequently, we construct a Mandarin SER dataset - SEQD, to demonstrate the effectiveness of our method in improving model performance. Finally, two experiments demonstrate the effectiveness of our method on both Mandarin and English SER tasks.

5. Acknowledgement

This paper is supported by the Key Research and Development Program of Guangdong Province under grant No.2021B0101400003. Corresponding author is Jianzong Wang from Ping An Technology (Shenzhen) Co., Ltd (jzwang@188.com).

6. References

- [1] R. Cowie, E. Douglas-Cowie, N. Tsapatsoulis, G. Votsis, S. Kollias, W. Fellenz, and J. G. Taylor, "Emotion recognition in human-computer interaction," *IEEE Signal processing magazine*, vol. 18, no. 1, pp. 32–80, 2001.
- [2] S. Si, J. Wang, J. Peng, and J. Xiao, "Towards speaker age estimation with label distribution learning," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 4618–4622.
- [3] S. Si, J. Wang, X. Qu, N. Cheng, W. Wei, X. Zhu, and J. Xiao, "Speech2video: Cross-modal distillation for speech to video generation," in *Proceedings of the Annual Conference of the International Speech Communication Association (INTERSPEECH)*, 2021.
- [4] D. Li, J. Liu, Z. Yang, L. Sun, and Z. Wang, "Speech emotion recognition using recurrent neural networks with directional self-attention," *Expert Systems with Applications*, vol. 173, p. 114683, 2021.
- [5] Z. Peng, Y. Lu, S. Pan, and Y. Liu, "Efficient speech emotion recognition using multi-scale cnn and attention," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 3020–3024.
- [6] W. Chen, X. Xing, X. Xu, J. Yang, and J. Pang, "Key-sparse transformer for multimodal speech emotion recognition," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6897–6901.
- [7] P. Singh, R. Srivastava, K. Rana, and V. Kumar, "A multimodal hierarchical approach to speech emotion recognition from audio and text," *Knowledge-Based Systems*, vol. 229, p. 107316, 2021.
- [8] S. Padi, S. O. Sadjadi, D. Manocha, and R. D. Sriram, "Multimodal emotion recognition using transfer learning from speaker recognition and bert-based models," *arXiv preprint arXiv:2202.08974*, 2022.
- [9] X. Cai, J. Yuan, R. Zheng, L. Huang, and K. Church, "Speech emotion recognition with multi-task learning," in *IEEE Conference of the International Speech Communication Association (INTERSPEECH)*, vol. 2021, 2021.
- [10] M. Crawshaw, "Multi-task learning with deep neural networks: A survey," *arXiv preprint arXiv:2009.09796*, 2020.
- [11] Y. Li, T. Zhao, and T. Kawahara, "Improved end-to-end speech emotion recognition using self attention mechanism and multitask learning," in *IEEE Conference of the International Speech Communication Association (INTERSPEECH)*, 2019, pp. 2803–2807.
- [12] A. Nediyanath, P. Paramasivam, and P. Yenigalla, "Multi-head attention for speech emotion recognition with auxiliary learning of gender recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7179–7183.
- [13] S. E. Eskimez, Z. Duan, and W. Heinzelman, "Unsupervised learning approach to feature analysis for automatic speech emotion recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5099–5103.
- [14] S. Si, J. Wang, H. Sun, J. Wu, C. Zhang, X. Qu, N. Cheng, L. Chen, and J. Xiao, "Variational information bottleneck for effective low-resource audio classification," in *Proceedings of the Annual Conference of the International Speech Communication Association (INTERSPEECH)*, 2021, p. 31.
- [15] 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, pp. 335–359, 2008.
- [16] S. R. Livingstone and F. A. Russo, "The ryerson audio-visual database of emotional speech and song (ravdess): A dynamic, multimodal set of facial and vocal expressions in north american english," *PLoS one*, vol. 13, no. 5, p. e0196391, 2018.
- [17] K. Zhou, B. Sisman, R. Liu, and H. Li, "Emotional voice conversion: Theory, databases and esd," *Speech Communication*, vol. 137, pp. 1–18, 2022.
- [18] H. Cao, D. G. Cooper, M. K. Keutmann, R. C. Gur, A. Nenkova, and R. Verma, "Crema-d: Crowd-sourced emotional multimodal actors dataset," *IEEE transactions on affective computing*, vol. 5, no. 4, pp. 377–390, 2014.
- [19] J. T. F. L. M. Zhang and H. Jia, "Design of speech corpus for mandarin text to speech," in *The Blizzard Challenge*, 2008.
- [20] R. Plutchik, "A general psychoevolutionary theory of emotion," in *Theories of emotion*. Elsevier, 1980, pp. 3–33.
- [21] E. Cambria, A. Livingstone, and A. Hussain, "The hourglass of emotions," in *Cognitive behavioural systems*. Springer, 2012, pp. 144–157.
- [22] B. Desplanques, J. Thienpondt, and K. Demuyck, "Ecapadnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification," *arXiv preprint arXiv:2005.07143*, 2020.
- [23] S.-H. Gao, M.-M. Cheng, K. Zhao, X.-Y. Zhang, M.-H. Yang, and P. Torr, "Res2net: A new multi-scale backbone architecture," *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 2, pp. 652–662, 2019.
- [24] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.
- [25] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 369–376.
- [26] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2980–2988.
- [27] I. Lawrence and K. Lin, "A concordance correlation coefficient to evaluate reproducibility," *Biometrics*, pp. 255–268, 1989.
- [28] B. McFee, C. Raffel, D. Liang, D. P. Ellis, M. McVicar, E. Battenberg, and O. Nieto, "librosa: Audio and music signal analysis in python," in *Proceedings of the 14th python in science conference*, vol. 8. Citeseer, 2015, pp. 18–25.
- [29] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, "Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline," in *20th Conference of the Oriental Chapter of the International Coordinating Committee on Speech Databases and Speech I/O Systems and Assessment (O-COCOSDA)*. IEEE, 2017, pp. 1–5.
- [30] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [31] J. Kominek and A. W. Black, "The cmu arctic speech databases," in *Fifth ISCA workshop on speech synthesis*, 2004.
- [32] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [33] Z. Zhao, Z. Bao, Z. Zhang, N. Cummins, H. Wang, and B. Schuller, "Attention-enhanced connectionist temporal classification for discrete speech emotion recognition," *Proceedings of the IEEE International Speech Communication Association (INTERSPEECH)*, 2019.
- [34] Y. Xu, H. Xu, and J. Zou, "Hgm: A hierarchical grained and feature model for acoustic emotion recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6499–6503.
- [35] M. Xu, F. Zhang, and W. Zhang, "Head fusion: improving the accuracy and robustness of speech emotion recognition on the iemocap and ravdess dataset," *IEEE Access*, vol. 9, pp. 74 539–74 549, 2021.