

An Efficient Transfer Learning Method Based on Adapter with Local Attributes for Speech Emotion Recognition

Haoyu Song* and Ian McLoughlin* and Qing Gu† and Nan Jiang† and Yan Song†

* Singapore Institute of Technology, Singapore

E-mail: 2303822@sit.singaporetech.edu.sg, ian.mcloughlin@singaporetech.edu.sg

Tel: +65 6592 1189

† School of Information Science and Technology, University of Science and Technology of China, Hefei, China

E-mail: qinggu6@mail.ustc.edu.cn, jiang_nan@mail.ustc.edu.cn, songy@ustc.edu.cn

Abstract—Existing speech emotion recognition (SER) methods commonly suffer from the lack of high-quality large-scale corpus, partly due to the complex, psychological nature of emotion which makes accurate labeling difficult and time consuming. Recently, transfer learning based methods that exploit the encoders pretrained on large-scale speech corpus (*e.g.*, Wav2Vec2.0 and HuBERT) have shown strong potential for downstream SER tasks. However, task-specific finetuning remains necessary for various conversational scenarios of different topics, speakers and languages to achieve satisfactory performance. It generally requires costly encoder retraining for individual SER tasks. To address this issue, we propose to train an adapter with local attributes for efficient transfer learning. Specifically, a weighted average pooling-Transformer (WAP-Transformer) is proposed as a lightweight backbone to enrich the frame-level representation. An adapter with teacher-student branches is exploited for task-agnostic transfer learning, where the student branch is jointly optimized via mask prediction and self-distillation objectives, and the teacher branch is obtained online from the student via exponential moving average (EMA). Meanwhile, local attributes are learned from the teacher branch via unsupervised clustering, which aims to act as a universal model that provides additional semantic-rich supervisions. A statistical attentive pooling (SAP) module is proposed to obtain utterance representation for finetuning. To evaluate the effectiveness of the proposed adapter with local attributes, extensive experiments have been conducted on IEMOCAP. Superior performance has been reported, compared to the previous state-of-the-art methods in similar settings.

I. INTRODUCTION

Speech Emotion Recognition (SER) has gained considerable attention due to its potential in application scenarios like human-computer interaction, affective computing, and mental health monitoring [1], [2]. The task remains challenging, particularly in cross-session settings, where models must generalize to unseen speakers with diverse styles, prosodies, dialects, languages, and acoustic variations [3], [4], [5].

Unlike automatic speech recognition (ASR), where phonetic content is relatively stable across different individuals [6], emotional expressions are inherently subjective and influenced by cultural background, speaker physiology, and recording conditions [7]. It is thus necessary to collect a large-scale high-quality emotion corpus with wide coverage of speakers,

languages and topics to guarantee the SER performance. However, due to the expensive annotation cost, existing SER methods [8], [9], [10], [11], [12] often struggle with few-shot learning and domain shift issues.

Recently, transfer learning based SER methods, that employ pretrained Wav2Vec2.0 and HuBERT [13], [14] as front-end encoders, have demonstrated the ability to address data scarcity issues [15], [16], [17]. However, it is still necessary to retrain the encoder for individual SER tasks to achieve satisfying performance – something computationally complex and requiring a large memory space. In [18], Ren et al. presented a self-distillation method to simultaneously fine-tune a pretrained Wav2vec2.0 model and train a shallow version of it for fast yet effective SER. In [3], a local prototypical mapping network (LPMN) was proposed to characterize the complex distribution of latent embedding spaces via an unsupervised learning method similar to VQ-VAE [19]. A prototype selection scheme is used to reduce bias caused by irrelevant factors during post-processing for specific SER tasks.

In this paper, we propose to learn an adapter with local attributes for various SER tasks (*e.g.*, that are cross-session validated using the Interactive Emotional Dyadic Motion Capture (IEMOCAP) corpus [20]). As shown in Fig. 1, a weighted average pooling Transformer (WAP-Transformer) is first proposed as a lightweight backbone to provide enriched frame-level representation. An adapter with teacher-student branches is then designed for self-supervised transfer learning. The student branch is optimized with a mask prediction objective similar to [21], [22], [23], from which the teacher branch are online updated via exponential moving average (EMA). Meanwhile, we perform self-distillation on masked frames and take the teacher network as an online tokenizer, from which the local attributes are simultaneously learned to characterize the frame-level feature distribution via batch-wise unsupervised clustering. In effect, the learned local attributes act as the universal model to enhance the task-agnostic adapter training for various downstream SER tasks, partly addressing the domain shift issues. Furthermore, a statistical attentive pooling (SAP) is proposed to aggregate the output of the adapter into an

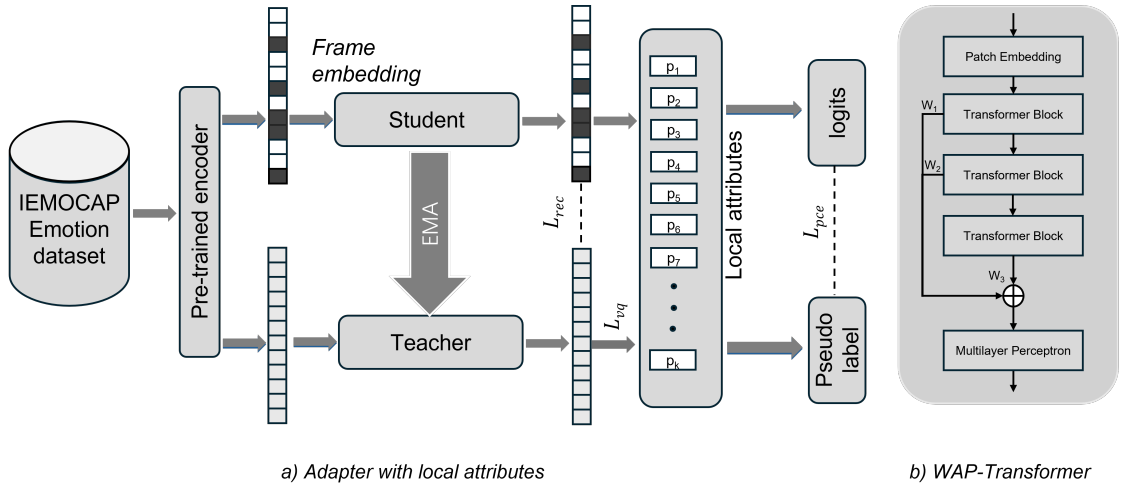


Fig. 1: a) The architecture of adapter training with mask prediction and online self-distillation on extracted frame-level feature space, b) The structure of WAP-Transformer used as backbone of adapter

utterance-level representation, which is classified according to the specific SER task. It is worth noting that only the lightweight adapter and classifier needs to be finetuned, which is more efficient than retraining the while encoder.

We perform 5-fold cross-validation with cross-session settings for evaluation [3], [24]. Experimental results on the IEMOCAP [20] benchmark demonstrate superior performance, achieving 78.32%, 77.56% and 75.63% unweighted accuracy (UA), weighted accuracy (WA) and F1-score, respectively.

II. METHODOLOGY

A. Overview of Adapter with Local Attributes

The architecture of our proposed adapter with local attributes is shown in Fig. 1, which consists of: 1) a pretrained encoder *i.e.*, HuBERT-large-960 [13], 2) a lightweight WAP-Transformer, and 3) the teacher-student branches.

1) **Pretrained encoder**: This consists of a 7-layer convolutional block, designed to capture low-level spectral and temporal cues, followed by a deep 24-layer Transformer-based context network with a hidden dimensions of 1024, 16 attention heads, and an intermediate feed-forward dimension of 4096. Pretrained on 960 hours of speech data, this model yields robust and general acoustic representations.

2) **Lightweight WAP-Transformer**: This consists of 3 transformer layers, with the outputs of each layer being adaptively pooled to enrich the frame-level representation. The aim is to reduce the adapter model size. The WAP-Transformer input is a sequence denoted by $X = X_{emb} + E_{pos}$, where the embeddings $X_{emb} \in R^{T \times D}$ are obtained by feeding the input waveform to the pretrained encoder. Following [22], a standard learnable 1D position embedding $E_{pos} \in R^{T \times D}$ is added to frame-level, and T is sequence length. The output of the l -th transformer layer is conceptually defined as $X^l = \text{Transformer}(X^{l-1})$. The output of WAP-transformer Z is

weighted averaged $Z = \sum_{l=1}^L w_l X^l$ with learnable parameters $W = [w_1, \dots, w_L]$, where $L = 3$,

3) **Teacher-student branches**: A Siamese structure with lightweight WAP-Transformer as backbone is exploited to perform self-supervised transfer learning on a given emotion corpus. As aforementioned, the student branch is optimized with both mask prediction and self-distillation objectives, and the teacher is ensembled from the students over epochs. The training process is detailed as follows, consisting of mask prediction, local attribute learning and frame-level self distillation.

B. Adapter training

1) **Mask Prediction**: The input to the student branch X' is the masked version of sequence X , where 40% of positions M are randomly selected and replaced with learnable embedding $e_{[M]} \in R^D$. By feeding X' to the student branch *i.e.*, WAP-Transformer E_{θ}^s , the reconstructed sequence $Z^s = \{z_i^s\}_{i=1}^T$ is obtained by predicting from the remaining visible ones according to context information. The input to the teacher branch E_{θ}^t is X , and the sequence $Z^t = \{z_i^t\}_{i=1}^T$ is used as the target of the masked prediction for optimizing the student parameters θ^s . The mask reconstruction loss is formulated as the mean square error of masked positions, which is

$$\mathcal{L}_{\text{rec}} = \frac{1}{\|M\|} \sum_{i \in M} \|z_i^t - z_i^s\|_2^2 \quad (1)$$

And the parameter of teacher θ^s is updated using EMA, that is

$$\theta_t^{i+1} = \alpha \theta_t^i + (1 - \alpha) \theta_s^{i+1} \quad (2)$$

where i denotes the learning iteration, and α is the smoothing hyperparameter, set to 0.999 by default. The large α is used to smooth the rapidly changing student parameters. This makes the teacher model more stable, avoids overfitting to noise, and provides a consistent signal that improves convergence and generalization.

This EMA smooths out the short-term noise in the student’s rapidly changing parameters, producing a more stable teacher signal that accelerates convergence and improves generalization.

In [22], [13], an offline additional tokenizer is further introduced to represent the utterance as a sequence of discrete tokens. We propose to learn the local-attribute prototypes directly from the teacher branch; these prototypes offer semantically rich supervision that complements the masked-prediction objective. Because each prototype is the quantiser output, its index $k(z^t)$ naturally serves as the pseudo-label for frame-level self-distillation, while the student produces its logits by computing the (negative) distances between its own embedding and all prototypes (see next subsection).

2) **Local Attributes Learning:** Local attributes $P = \{p_1, p_2, \dots, p_K\}$ are learned to characterize the frame-level feature distribution, which can be learned by a batch-wise unsupervised clustering method. Given a batch of the extracted frame-level embeddings Z^t from the teacher branch, local attributes P can be learned as follows

$$\mathcal{L}_{VQ} = \sum_{k=1}^K \sum_{z^t \in p_k} \|z^t - p_k\|_2^2 \quad (3)$$

where $z^t \in p_i$ denotes embeddings nearest to the i -th attribute p_i . That is, we can assign z^t to pseudo-label i , according to Euclidean distance

$$i = \arg \min_k \|z^t - p_k\|_2^2, \quad (4)$$

And the corresponding attribute is updated as

$$p'_i \leftarrow p_i + \eta_t(z^t - p_i) \quad (5)$$

The learned attributes may be considered as a universal model to characterize the distribution of frame-level embeddings. We further perform self-distillation to update the student branch with local attributes, besides the reconstruction loss L_{rec} in eqn.(1).

3) **Frame-level Self-Distillation:** We perform frame-level self-distillation on the masked positions M . From the teacher branch, we can obtain the pseudo-labels of Z^t according to eqn. (4). For the student branch, the logits can be obtained by computing the cosine similarity between of embeddings Z^s and local attributes P . The pseudo cross entropy (CE) loss \mathcal{L}_{pce} can then be computed to jointly optimize with \mathcal{L}_{rec} , that is

$$\mathcal{L} = \lambda \mathcal{L}_{rec} + (1 - \lambda) \mathcal{L}_{pce} \quad (6)$$

where λ is the hyperparameter for balance reconstruction and pseudo CE loss. A ramp-down scheme from 1.0 to 0.5 is applied to place more emphasis on \mathcal{L}_{rec} at the early stage, allowing the student to first learn stable representations by mimicking the teacher. As training progresses, the weight gradually shifts toward \mathcal{L}_{pce} , enabling the model to focus more on discriminative learning. The lower bound of 0.5 ensures that \mathcal{L}_{rec} still contributes throughout training, preventing the student from completely drifting away from the teacher guidance.

C. Finetune on SER task

We follow [24], [3] to further finetune the adapter for SER. An SAP module comprising a learnable convolution and a statistic pooling layer is proposed to exploit frame-level attention. The convolution layer weighs the frame-level features, while a statistic pooling layer collects the 1st- and 2nd-order weighted statistics. Specifically, let $Z \in R^{T \times D}$ be the output features for L frames with D dimensions. By feeding Z into a 1×1 convolution layer, followed by softmax activation along the time axes, an attentive weight map $A \in R^T \times K$ can be obtained, with K attention heads. Statistical pooling is performed:

$$\begin{aligned} \mu &= Z^T A \\ \sigma^2 &= Z^T (A \odot A) - (Z^T A \odot Z^T A) \end{aligned} \quad (7)$$

where \odot denotes the Hadamard product. The μ and σ^2 are L_2 -normalized and concatenated as the input the final emotion classification layer. IEMOCAP 5-fold cross-validation is used for fine-tuning, where the dataset is divided into five sessions (each with 2 speakers, for a total of 10 speakers). In each fold, one session serves as the validation set while the remaining four sessions are used for training.

III. EXPERIMENTS

A. Dataset and Evaluation Metrics

IEMOCAP contains recordings of multiple actors in improvised and scripted dialogues with a wide variation in speaker and session characteristics [20]. It is widely used for SER research and allows us to compare the effectiveness of adapter to other systems. We assessed performance using WA, UA, and macro F1 scores.


WA computes a weighted average of per-class accuracy, while UA, as a simple average, better reflects fairness in the presence of class imbalance. The macro F1-score, computed as the average F1 across all classes, provides a balanced evaluation by giving equal weight to each class, which is particularly useful under class imbalance.

The results, presented below, confirm the effectiveness of proposed data augmentation strategy in enhancing robustness.


B. Implementation details

As aforementioned, the adapter module exploits the WAP-Transformer backbone, acting on the output of the pretrained encoder (HuBERT-large-960), as shown in Fig. 1. This begins with a patch embedding layer that projects 1024-dimensional inputs to 384 dimensions, followed by three Transformer blocks at the 384-dimensional level. A fourth block is a 3-layer MLP then applied for aggregation. A Siamese network based WAP-Transformer is designed as adapter, and corresponding K local attributes are learned as a universal model for self-distillation. A SAP module with 4 attention heads is used to project a sequence of frame-level embeddings into utterance-level representation for SER.

Adapter training is conducted with a batch size of 96 over 100 epochs using the Adam optimizer with an initial learning



(a)



(b)

Fig. 2: t-SNE projections emotion embeddings on unseen test data. (a) Emotion embeddings obtained from pretrained encoder, (b) Emotion embeddings obtained from our proposed adapter.

rate of $1e-4$ and a cosine learning rate scheduler. Fine-tuning also employs a batch size of 96 over 100 epochs, but uses cross-entropy loss for the primary emotion classification task with improvised dialogues, over five classes (with happy and excited merged to account for limited data). To further address data imbalance and leverage the reconstruction ability of our trained model, we apply data augmentation during fine-tuning by randomly masking parts of the embeddings for minority classes, thereby enhancing sample diversity and balancing class distributions.

C. Comparison with state-of-the-art methods

Table I presents 5-fold cross validation results in terms of UA, WA and F1-score, for adapter training and several state-of-the-art systems. Overall results demonstrate that our adapter can outperform existing methods on the challenging IEMOCAP dataset. We achieved a UA of 78.32% and an WA of 77.56%, surpassing previous approaches such as Co-attention, W2v2-PT [25], Spk-norm [26], GLRF [27], and LPMN [3] in similar settings. For example, The competitive LPMN reported UA and WA values of 76.85% and 74.50% respectively. After a prototype selection for post processing, it can improve to 77.42% and 75.82%. Our proposed WAP-Transformer improves over LPMN-1024-s60 around 0.80% and 2.26% in terms of UA and WA metrics. These gains are particularly impressive given the inherent challenges of IEMOCAP, which include significant cross-session and cross-speaker variability that typically degrade performance. The results validate the effectiveness of combining a large-scaled pre-trained encoder, WAP-Transformer based adapter training, and finetuning.

Fig. 2 compares the utterance-level embedding spaces obtained from the same HuBERT encoder before and after frame-level self-distillation. (a) Without self-distillation, the clusters corresponding to the four emotions are loosely formed

TABLE I: Validation scores of the proposed WAP-Transformer compared with state-of-the-art methods on IEMOCAP.

Method	UA	WA	F1
Co-att [25]	72.70	71.64	–
W2v2-PT [15]	–	67.20	–
Spk-norm [26]	–	74.20	–
GLRF [27]	73.31	72.81	72.92
LPMN-1024-s60 [3]	77.42	75.82	75.71
LPMN [3]	76.85	74.50	74.88
WAP-Transformer (Ours)	78.32	77.56	75.63

and partially overlapped – particularly between Neutral and Happy. (b) Introducing frame-level self-distillation leads to markedly tighter intra-class cohesion and larger inter-class margins, yielding a much cleaner separation of Sad and Angry from the other emotions while also reducing the number of outliers. These observations confirm that self-distillation effectively drives the encoder to focus on affect-relevant acoustic cues, thereby enhancing the discriminability of the learned representations for downstream SER.

D. Ablation study to explore key components

To evaluate the impact of the key contributions in WAP-Transformer, we conduct an ablation study focusing on three critical factors: (1) the effect of adapter training via WAP-Transformer, (2) the influence of local attribute (LA) size in the VQ module Table II summarizes the performance under different configurations.

Overall, these ablation experiments reveal some useful findings. Firstly, adapter training can significantly boost performance with WAP-Transformer. Without adapter training, the model achieved a UA of 75.73%, a WA of 74.91%, and an F1 score of 72.90%. All metrics increased after adapter training, to 78.32% (UA), 77.56% (WA), and 75.63% (F1) respectively – around 2% absolute increase for each. This confirms that adapter training benefits from WAP-Transformer to learn local

TABLE II: Ablation studies on the impact of WAP-LA components and training settings.

<i>K</i> attributes	UA (%)	WA (%)	F1-score
<i>K</i> = 1024	78.32	77.56	75.63
<i>K</i> = 512	77.22	76.93	74.26
<i>K</i> = 256	77.00	76.21	75.14
<i>K</i> = 128	76.10	76.28	73.80
WAP-Transformer (w/o Adapter Training)	75.73	74.91	72.90

attributes, capturing semantic-rich information like speakers, phonetic content and styles.

When we consider local attributes (LA), Table II indicates that a size of 1024 yields the best performance of 78.32% and 77.56% and 75.63% in UA, WA and F1-score. Reducing the LA size degrades performance, suggesting that a larger size is more effective at capturing fine-grained details, although it may also increase computational cost and the risk of overfitting.

Ablation studies collectively confirm that the component of the proposed approach – adapter training, an appropriate LA size, and careful integration of domain information all play a crucial role in improving performance.

IV. CONCLUSION

This paper has proposed WAP-Transformer to effectively reduce the gap between a large-scale pretrained HuBERT model and downstream SER tasks with limited corpus size. WAP-Transformer introduces a lightweight transformer (*i.e.*, WAP-Transformer) based adapter, which contains a teacher and student branch. WAP-Transformer learns local attributes that can effectively capture the semantic-rich information related to speakers, content and styles for different tasks. After adapter training and adjustment, we finetune the adapter to perform SER tasks via an attentive pooling layer and a classification layer. 5-fold cross validation on IEMOCAP demonstrates superior performance of 78.32%, 77.56% and 75.63% for UA, WA and F1-scores respectively. Ablation studies confirm the effectiveness of the main contributions introduced by WAP-Transformer.

REFERENCES

- [1] B. Schuller, A. Batliner, S. Steidl, and D. Seppi, “Recognising realistic emotions and affect in speech: State of the art and lessons learnt from the first challenge,” *Speech Communication*, vol. 53, no. 9-10, pp. 1062–1087, 2011.
- [2] T. Reitmaier, E. Wallington, D. Kalarikalayil Raju, O. Klejch, J. Pearson, M. Jones, P. Bell, and S. Robinson, “Opportunities and challenges of automatic speech recognition systems for low-resource language speakers,” in *Proc. of CHI*, pp. 1–17, 2022.
- [3] Y. Xi, Y. Song, L. Dai, H. Song, and I. McLoughlin, “An effective local prototypical mapping network for speech emotion recognition,” in *Proc. of Interspeech*, pp. 1055–1059, 2024.
- [4] Y. Yin, B. Huang, Y. Wu, and M. Soleymani, “Speaker-invariant adversarial domain adaptation for emotion recognition,” in *Proc. of ICMI 2020*, pp. 481–490, 2020.
- [5] S. Li, P. Song, L. Ji, Y. Jin, and W. Zheng, “A generalized sub-space distribution adaptation framework for cross-corpus speech emotion recognition,” in *Proc. of ICASSP*, 2023.

- [6] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, *et al.*, “Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups,” *IEEE Signal processing magazine*, vol. 29, no. 6, pp. 82–97, 2012.
- [7] D. Ververidis and C. Kotropoulos, “Emotional speech recognition: Resources, features, and methods,” *Speech communication*, vol. 48, no. 9, pp. 1162–1181, 2006.
- [8] Y. Yu and Y.-J. Kim, “Attention-lstm-attention model for speech emotion recognition and analysis of iemocap database,” *Electronics*, vol. 9, no. 5, p. 713, 2020.
- [9] C.-S. Ahn, C. Kasun, S. Sivadas, and J. Rajapakse, “Recurrent multi-head attention fusion network for combining audio and text for speech emotion recognition,” in *Proc. Interspeech*, pp. 744–748, 2022.
- [10] H. Zou, Y. Si, C. Chen, D. Rajan, and E. S. Chng, “Speech emotion recognition with co-attention based multi-level acoustic information,” in *Proc. of ICASSP*, pp. 7367–7371, 2022.
- [11] C. Fu, C. Liu, C. T. Ishi, and H. Ishiguro, “Maec: Multi-instance learning with an adversarial auto-encoder-based classifier for speech emotion recognition,” in *Proc. of ICASSP*, pp. 6299–6303, 2021.
- [12] S. Mao, P. Ching, and T. Lee, “Deep Learning of Segment-Level Feature Representation with Multiple Instance Learning for Utterance-Level Speech Emotion Recognition,” in *Proc. of Interspeech*, pp. 1686–1690, 2019.
- [13] W. N. Hsu, B. Bolte, Y. H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” pp. 3451–3460, 2021.
- [14] A. Baeviski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in neural information processing systems*, vol. 33, pp. 12449–12460, 2020.
- [15] L. Pepino, P. Riera, and L. Ferrer, “Emotion recognition from speech using wav2vec 2.0 embeddings,” *arXiv preprint arXiv:2104.03502*, 2021.
- [16] L. Zhang, B. Schuller, and M. Chen, “Multimodal emotion recognition: Challenges and opportunities,” in *Proc. of ICASSP*, pp. 8094–8098, 2019.
- [17] A. Geetha, T. Mala, D. Priyanka, and E. Uma, “Multimodal emotion recognition with deep learning: advancements, challenges, and future directions,” *Information Fusion*, vol. 105, p. 102218, 2024.
- [18] Z. Ren, T. T. Nguyen, Y. Chang, and B. W. Schuller, “Fast yet effective speech emotion recognition with self-distillation,” in *Proc. of ICASSP*, pp. 1–5, 2023.
- [19] A. van den Oord, O. Vinyals, and K. Kavukcuoglu, “Neural discrete representation learning,” *CoRR*, vol. abs/1711.00937, 2017.
- [20] Busso, Carlos, Bulut, Murtaza, Lee, C. Chun, Kazemzadeh, Abe, Mower, Emily, Kim, Samuel, Chang, J. N., Lee, Sungbok, Narayanan, and S. S., “IEMOCAP: interactive emotional dyadic motion capture database,” *Language Resources and Evaluation*, vol. 42, no. 4, p. 335, 2008.
- [21] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, “Masked autoencoders are scalable vision learners,” in *Proc. of the CVPR*, pp. 16000–16009, 2022.
- [22] H. Bao, L. Dong, S. Piao, and F. Wei, “Beit: Bert pre-training of image transformers,” *arXiv preprint arXiv:2106.08254*, 2021.
- [23] Z. Xie, Z. Zhang, Y. Cao, Y. Lin, J. Bao, Z. Yao, Q. Dai, and H. Hu, “SimMIM: A simple framework for masked image modeling,” in *Proc. of CVPR*, pp. 9653–9663, June 2022.
- [24] Y.-X. Xi, Y. Song, L.-R. Dai, I. McLoughlin, and L. Liu, “Frontend attributes disentanglement for speech emotion recognition,” in *Proc. of ICASSP*, pp. 7712–7716, 2022.
- [25] H. Zou, Y. Si, C. Chen, D. Rajan, and E. S. Chng, “Speech emotion recognition with co-attention based multi-level acoustic information,” in *Proc. of ICASSP*, pp. 7367–7371, 2022.
- [26] I. Gat, H. Aronowitz, W. Zhu, E. Morais, and R. Hoory, “Speaker normalization for self-supervised speech emotion recognition,” in *Proc. of ICASSP*, pp. 7342–7346, 2022.
- [27] C. Ding, J. Li, D. Zong, B. Li, T. Zhang, and Q. Zhou, “Stable speech emotion recognition with head-k-pooling loss,” *INTERSPEECH*, 2023.