

# IMPROVING SPEECH EMOTION RECOGNITION WITH UNSUPERVISED SPEAKING STYLE TRANSFER

Leyuan Qu<sup>1</sup>, Wei Wang<sup>2</sup>, Cornelius Weber<sup>3</sup>, Pengcheng Yue<sup>1</sup>, Taihao Li<sup>1\*</sup>, Stefan Wermter<sup>3</sup>

<sup>1</sup>Zhejiang Lab, China, <sup>2</sup>Xinjiang University, China, <sup>3</sup>University of Hamburg, Germany

Humans can effortlessly modify various prosodic attributes, such as the placement of stress and the intensity of sentiment, to convey a specific emotion while maintaining consistent linguistic content. Motivated by this capability, we propose EmoAug, a novel style transfer model designed to enhance emotional expression and tackle the data scarcity issue in speech emotion recognition tasks. EmoAug consists of a semantic encoder and a paralinguistic encoder that represent verbal and non-verbal information respectively. Additionally, a decoder reconstructs speech signals by conditioning on the aforementioned two information flows in an unsupervised fashion. Once training is completed, EmoAug enriches expressions of emotional speech with different prosodic attributes, such as stress, rhythm and intensity, by feeding different styles into the paralinguistic encoder. EmoAug enables us to generate similar numbers of samples for each class to tackle the data imbalance issue as well. Experimental results on the IEMOCAP dataset demonstrate that EmoAug can successfully transfer different speaking styles while retaining the speaker identity and semantic content. Furthermore, we train a SER model with data augmented by EmoAug and show that the augmented model not only surpasses the state-of-the-art supervised and self-supervised methods but also overcomes overfitting problems caused by data imbalance. Some audio samples can be found on our demo website<sup>1</sup>.

**Index Terms**— Speech emotion recognition, data augmentation, style transfer.

## 1. INTRODUCTION

Speech Emotion Recognition (SER) aims at recognizing and understanding human emotions from spoken language, which can significantly benefit and promote the experience of human-machine interaction. Plenty of studies have investigated SER in recent years [1]. However, the performance of SER is constrained by the lack of large-scale labelled datasets. Most speech emotion datasets are collected under simulated or elicited scenarios, since capturing natural and spontaneous emotional speech is very challenging. Furthermore, different people perceive emotions differently, which may result in

ambiguity in data labeling, especially in emotions with weak intensity [2]. Moreover, data imbalance of different emotions is another problem that can lead to model overfitting to some frequent emotions, such as “Neutral”.

Data augmentation is treated as a promising method to address these issues. Most of the augmentation methods used in SER are borrowed from Automatic Speech Recognition (ASR), such as SpecAugment [3], Vocal Tract Length Perturbation (VTLP) [4], speed perturbation [5], pitch shift or noise injection. However, while the variations in pitch or speed do not change the semantic content, they may have an effect on emotion expressions. For instance, sad emotions are often conveyed at slow speed while angry emotions tend to be expressed fast.

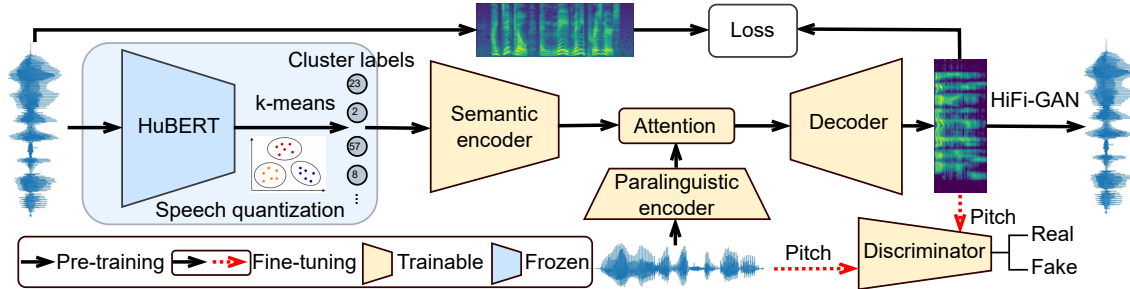
Alternatively, different variants of models are adopted to generate intermediate emotional features or alter carried emotions while keeping speech content unchanged, for instance, star Generative Adversarial Networks (GANs) [6], CycleGAN [7] and global style token [8]. However, the generated intermediate features are not easy to evaluate intuitively. In addition, prosody expression is strongly associated with speech content, and the altered emotions by GANs may cause conflicts or ambiguity between speech prosody and content.

Humans can easily alter different prosody attributes, such as stress position and sentiment intensity, to express a given emotion with invariant linguistic content [9]. For example, when expressing sad emotions with the semantic content “I am not happy today”, one can emphasize “not happy” or put stress on “today”. Inspired by this capability, we propose EmoAug to vary prosody attributes and augment emotional speech while keeping emotions unchanged. EmoAug is trained in an unsupervised manner, which requires neither paired speech nor emotion labels. The main contributions of this paper are as follows:

1. We propose a novel unsupervised speaking style transfer model which enriches emotion expressions by altering stress, rhythm and intensity while keeping emotions, semantics and speaker identity invariant.
2. We implement quantized units to represent semantic speech content instead of using a well-trained ASR model, which enables us to work on emotional audio without text transcriptions.

\*Corresponding author (email: lith@zhejianglab.com)

<sup>1</sup><https://leyuanqu.github.io/EmoAug/>



**Fig. 1:** Overview of EmoAug. The model comprises two encoders and an attention-based decoder. In the pre-training phase, the decoder reconstructs input mel-spectrograms using representations acquired from the semantic and paralinguistic encoders. The loss function is computed by measuring the Mean Square Error (MSE) between the generated and original mel-spectrograms. During fine-tuning, style transfer is performed by directly substituting the input of the paralinguistic encoder with a target speaking style reference. Additionally, to enhance the quality of the converted audio, a discriminator is employed to differentiate between real and generated pitch contours.

- SER models trained with EmoAug outperform the state-of-the-art models by a large margin, effectively overcoming overfitting issues caused by data scarcity and class imbalance.

## 2. METHOD

An overview of EmoAug is shown in Fig. 1, which consists of a speech quantization module, a semantic encoder, a paralinguistic encoder and an attention-based decoder. An additional discriminator is utilized during the fine-tuning phase.

### 2.1. Speech Quantization and Semantic Encoder

Acquiring semantic speech information typically relies on a proficient ASR system. Nevertheless, training a reliable ASR model proves challenging due to the extensive labeled data required. In addition, ASR models are not robust to unseen noise or emotional speech [10]. Therefore, we employ HuBERT representations [11] to capture semantic speech content. These representations are acquired through self-supervised training, eliminating the need for human annotations in ASR training. HuBERT empowers us to address emotional speech and diverse styles that might substantially compromise the performance of an ASR model. As shown in Eq.1, the input speech signal  $x = (x_0, \dots, x_t)$  is firstly embedded into continuous vectors by the pre-trained HuBERT<sup>2</sup>, followed by the k-means algorithm that quantizes the continuous speech representations into discrete cluster labels  $u = (u_0, \dots, u_t)$ , e.g. “23, 23, 2, 2, ..., 57”.

$$u = k\text{-means}(\text{HuBERT}(x)) \quad (1)$$

We investigated how different vocabulary sizes of HuBERT impact in model performance. In this study, we employed a vocabulary size of 200 clusters, which yielded considerably improved performance compared to using 50 or 100 classes. To refine the semantic content, we proceed to eliminate repetitions and filter out tempo information ( $u \rightarrow \tilde{u}$ ), e.g. “23, 23, 2, 2, ..., 57”  $\rightarrow$  “23, 2, ..., 57”. After quantization, we map the cluster labels into latent representations with

the semantic encoder (*Sem*) which is comprised of three 512-channel Conv1D layers with kernel width of 5 and padding size of 2, and one bidirectional Long Short-Term Memory (LSTM) with 256 dimensions to capture local and global context information, respectively.

### 2.2. Paralinguistic Encoder

The paralinguistic encoder (*Par*) aims to learn utterance-level non-verbal information from input audio  $x$ , which contains speaking styles, emotion states, speaker identities, and so on. It is based on the ECAPA-TDNN model [12] which has been proposed for the task of speaker verification. The model begins with one Conv1D layer, a ReLU function, and Batch Normalization (BN), followed by three SE-Res2 blocks. Residual connections between the SE-Res2 blocks deliver different level outputs to the feature aggregation layer (Conv1D+ReLU). Subsequently, the aggregated outputs are dynamically weighed by the attentive statistics pooling layer (Conv1D+Tanh+Conv1D+Softmax), and then mapped to a fixed dimension by the last Fully Connected (FC) layer. We initialize ECAPA-TDNN with the pre-trained model<sup>3</sup>. The attentive statistics pooling, fixed dimension mapping and weight initialization prevent semantic information to leak from the paralinguistic encoder.

### 2.3. Decoder

Our decoder is based on the Tacotron2 [13] which uses location-aware attention (*Att*) [14] to connect the encoders and the decoder (see Eq. 2).

$$\tilde{m} = \text{Dec}(\text{Att}(\text{Sem}(\tilde{u}), \text{Par}(x))) \quad (2)$$

The decoder (*Dec*) generates one frame per time step in an auto-regressive fashion. Two FC layers map the ground-truth mel-spectrograms  $x$  into latent representations that are used by an LSTM module for teacher-forcing training. Finally, a FC layer maps the intermediate features to the dimension of input mel-spectrograms. The generated mel-spectrograms  $\tilde{m}$  are then transformed to waveforms  $\tilde{x}$  by HiFiGAN [15].

<sup>2</sup><https://huggingface.co/facebook/hubert-base-ls960>

<sup>3</sup><https://huggingface.co/speechbrain/spkrec-ecapa-voxceleb>

$$\tilde{x} = HiFiGAN(\tilde{m}) \quad (3)$$

## 2.4. Discriminator

After pre-training, we observed when feeding a different reference audio to the paralinguistic encoder for style transfer, the generated speech signals exhibit some distortions. To enhance the quality of the generated speech, we implement a discriminator to differentiate between the genuine and synthesized speech. Importantly, the discriminator exclusively differentiates pitch changes while preserving semantic content. The discriminator starts with three convolutional blocks (Conv1D+ReLU+BN+dropout), followed by two linear projection layers.

## 2.5. Speaking Style Transfer

We denote  $X^{s,e}$  as all utterances with emotion  $e$  uttered by speaker  $s$ . After pre-training, as shown in Eq. 4, speaking style transfer can be achieved by directly replacing the paralinguistic encoder input  $x^{s,e}$  with a target speaking style  $y^{s,e}$ , where  $x^{s,e} \in X^{s,e}$ ,  $y^{s,e} \in Y^{s,e}$  and  $Y^{s,e} = X^{s,e} - \{x^{s,e}\}$ .

$$\tilde{m}^{s,e} = Dec(Att(Sem(\tilde{u}), Par(y^{s,e}))) \quad (4)$$

Consequently, the converted mel-spectrograms  $\tilde{m}^{s,e}$  retain the same speaker identity, semantic content and emotions as the original audio  $x^{s,e}$ , but deliver different rhythms or intensities transferred from  $y^{s,e}$ . In addition, we generate different numbers of samples for each emotion to counter data imbalance.

## 3. EXPERIMENTAL SETUPS

### 3.1. Datasets

We utilize the LRS3-TED and IEMOCAP datasets for pre-training and fine-tuning respectively.

- **LRS3-TED** [16] is comprised of over 400 hours of video by more than 5000 speakers from TED and TEDx talks with spontaneous speech. LRS3-TED is collected with various speaking styles and emotions in a variety of acoustic scenes, which will help the model learn rich paralinguistic changes.
- **IEMOCAP** [17] is a multimodal emotion dataset recorded by 10 actors in a fictitious scenario. We follow the settings used in previous work [18, 19, 20], in which leave-one-session-out is adopted with 5-fold cross-validation. One session is left for testing and another one is used for validation, while the rest of the three sessions is utilized for training in each round. Four types of emotions (*happy*, *sad*, *angry* and *neutral*) with a total of 5531 utterances are considered.

### 3.2. Evaluation Metrics

We report the experimental results with Weighted Accuracy (WA) and Unweighted Accuracy (UA).

### 3.3. EmoAug Training

We first pre-train EmoAug on LRS3-TED and then fine-tune it on IEMOCAP for style transfer. During pre-training, the Adam optimizer is used with a weight decay rate of  $1e-6$ . The initial learning rate is  $1e-3$ , which is decayed with a factor of 0.9 after every 5000 iterations. In addition, gradient clipping with a threshold of 1.0, early stopping and scheduled sampling are adopted to avoid overfitting.

After pre-training, we fine-tune the model with a discriminator on the IEMOCAP dataset with small learning rates of  $1e-5$  and  $1e-4$  to improve the quality of generated speech.

### 3.4. SER Model Training

We perform SER by adding one additional FC layer on top of the HuBERT model which is pre-trained on large-scale unlabelled data by self-supervised learning. During training, different learning rates are used in the HuBERT model ( $1e-5$ ) and the FC layer ( $1e-4$ ) to retain the low-layer representations and enable the last FC layer to fit to the specific dataset.

## 4. EXPERIMENTAL RESULTS AND DISCUSSION

### 4.1. Comparison of Different Augmentation Times

We augment each utterance  $N$  times by transferring speaking styles from  $N$  randomly selected utterances that belong to the same speaker with the same type of emotion. We report the effect of different augmentation times and the discriminator on SER utilizing representations from HuBERT.

As can be seen in Fig. 2, considerable increases occur on UA along with the increase of augmentation time, where 0 means training SER models with only raw audio. Furthermore, the discriminator significantly enhances the quality of the generated speech, resulting in a substantial improvement in emotion recognition. The figure reveals that the speaking styles transferred by EmoAug effectively enrich the emotion expression on prosody and greatly enhance SER performance.


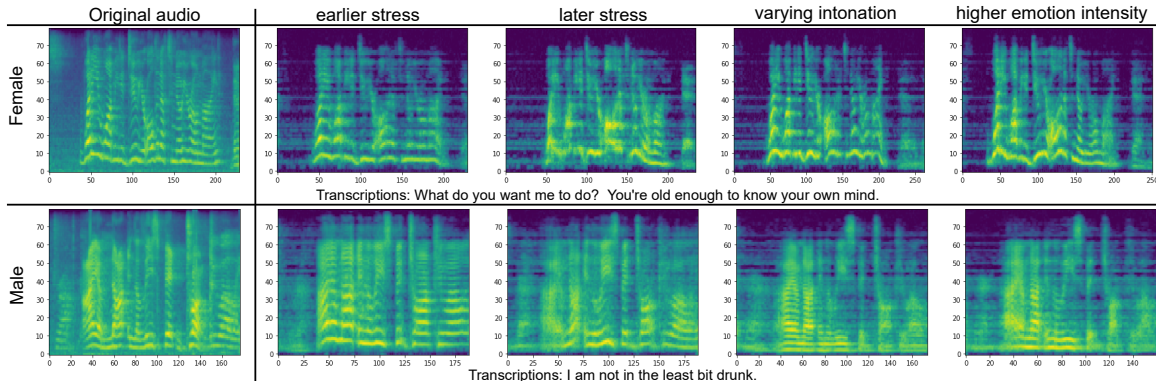


Fig. 2: The effect of different augmentation times on SER when training EmoAug with or without the discriminator.

In addition, we also visualize the original and generated mel-spectrograms. As shown in Fig. 3, in comparison to the original audio, EmoAug successfully transfers different speaking styles to the source audio while keeping the semantic information invariant. By listening to the generated audio, we found that EmoAug can effectively augment the

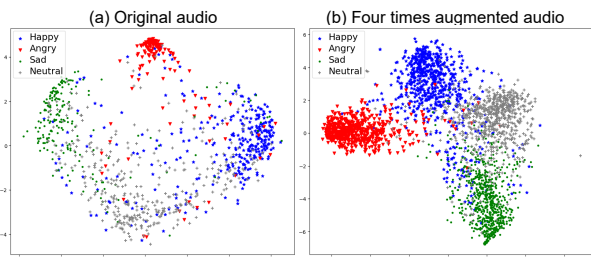


**Fig. 3:** A comparison of original and augmented mel-spectrograms with the *angry* emotion where the styles are transferred from the same speaker with the same emotion as the original audio. Different stress patterns, rhythms and intonations are shown in the augmented audio.

expression of source emotion by varying stress positions, intonations or even the intensity of emotions. We recommend readers to listen to the audio samples on our demo website.

#### 4.2. Model Performance with and without Augmentation

We select the utterances from session one (two speakers) of IEMOCAP and visualize the embeddings from the penultimate layer of the model trained with only original data and with four times augmented data. As depicted in Fig. 4 (a), the model trained only with original audio struggles to clearly separate the representations of neutral, sad and happy emotions. However, as shown in Fig. 4 (b), EmoAug successfully transfers speaking styles and enriches sample diversity, resulting in more precise and distinguishable model representations when trained with augmented audio.



**Fig. 4:** Comparison of models trained with original data (a) and augmented data (b), visualized with t-SNE.

#### 4.3. Comparison with Previous Work

We compare our methods with previous supervised and self-supervised models in Table 1. We also reimplement and augment IEMOCAP with emotional voice conversion models, CycleGAN [21] and StarGAN [6]. We reproduce the Copy-Paste [22] method by randomly concatenating two emotional utterances with the same emotion, which corresponds to the Same Emotion CopyPaste (SE-CP) setting in [22]. Additionally, we perturb speech on speed with the factors of 0.9, 1.0 and 1.1. Pitch shift is adopted by randomly raising or lowering 2 semitones on each audio. Table 1 shows that EmoAug outperforms previous methods by a big margin.

**Table 1:** Results of SER on IEMOCAP dataset with 5-fold cross-validation and leave-one-session-out settings.

Methods	WA	UA
<b>Supervised Methods</b>		
CNN-ELM+STC attention [18]	61.32	60.43
IS09-classification [19]	68.10	63.80
Co-attention-based fusion [20]	69.80	71.05
<b>Self-supervised Methods</b>		
Data2Vec Large [23]	66.31	-
WavLM Large [24]	70.62	-
HuBERT Large	70.24	71.13
<b>Emotional Voice Conversion Methods</b>		
HuBERT Large + CycleGAN [21]	71.57	72.02
HuBERT Large + StarGAN [6]	71.51	72.13
<b>Data Augmentation Methods</b>		
VTLP [25]	66.90	65.30
HuBERT Large + CopyPaste	70.79	71.35
HuBERT Large + Speed Perturbation	70.35	71.19
HuBERT Large + Pitch Shift	70.47	71.24
<b>Our Methods</b>		
HuBERT Large + EmoAug	<b>72.66</b>	<b>73.75</b>

## 5. CONCLUSION

We introduce EmoAug to tackle the challenges of data scarcity and data imbalance in SER. EmoAug is composed of a semantic encoder, a paralinguistic encoder, and a decoder to reconstruct speech in an unsupervised manner. Speaking style transfer is performed by altering the input of the paralinguistic encoder. Experimental results on the IEMOCAP dataset suggest that EmoAug can effectively enrich emotion expressions with different stress patterns, rhythms and intensities, and achieve superior performance in SER compared to state-of-the-art supervised and self-supervised methods.

## 6. ACKNOWLEDGEMENT

This work was supported in part by the National Science and Technology Major Project of China (2021ZD0114303), in part by the Youth Foundation Project of Zhejiang Lab (K2023KH0AA01), in part by the CML Project funded by the DFG, and in part by the Philosophy and Social Science Training Project of Xinjiang University (23CPY049).

## 7. REFERENCES

- [1] Johannes Wagner, Andreas Triantafyllopoulos, Hagen Wierstorf, Maximilian Schmitt, Felix Burkhardt, Florian Eyben, and Björn W Schuller, “Dawn of the transformer era in speech emotion recognition: closing the valence gap,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [2] Babak Joze Abbaschian, Daniel Sierra-Sosa, and Adel Elmaghraby, “Deep learning techniques for speech emotion recognition, from databases to models,” *Sensors*, vol. 21, no. 4, pp. 1249–1276, 2021.
- [3] Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and Quoc V Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *Proc. INTERSPEECH*, 2019, p. 1613–1617.
- [4] Navdeep Jaitly and Geoffrey E Hinton, “Vocal tract length perturbation (VTLP) improves speech recognition,” in *Proc. ICML Workshop on Deep Learning for Audio, Speech and Language*, 2013, vol. 117, pp. 21–25.
- [5] Tom Ko, Vijayaditya Peddinti, Daniel Povey, and Sanjeev Khudanpur, “Audio augmentation for speech recognition,” in *Proc. INTERSPEECH*, 2015, pp. 3586–3589.
- [6] Georgios Rizos, Alice Baird, Max Elliott, and Björn Schüller, “StarGAN for emotional speech conversion: Validated by data augmentation of end-to-end emotion recognition,” in *International Conference on Acoustics, Speech and Signal Processing*, 2020, pp. 3502–3506.
- [7] Fang Bao, Michael Neumann, and Ngoc Thang Vu, “CycleGAN-Based Emotion Style Transfer as Data Augmentation for Speech Emotion Recognition,” in *Proc. INTERSPEECH*, 2019, pp. 2828–2832.
- [8] Yuxuan Wang, Daisy Stanton, Yu Zhang, RJ-Skerry Ryan, Eric Battenberg, Joel Shor, Ying Xiao, Ye Jia, Fei Ren, and Rif A Saurous, “Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis,” in *ICML*. PMLR, 2018, pp. 5180–5189.
- [9] Björn Schüller and Anton Batliner, *Computational paralinguistics: emotion, affect and personality in speech and language processing*. John Wiley & Sons, 2013.
- [10] Egor Lakomkin, Mohammad Ali Zamani, Cornelius Weber, Sven Magg, and Stefan Wermter, “On the robustness of speech emotion recognition for human-robot interaction with deep neural networks,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 854–860.
- [11] Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed, “HuBERT: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [12] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuynck, “ECAPA-TDNN: emphasized channel attention, propagation and aggregation in TDNN based speaker verification,” in *Proc. INTERSPEECH*. 2020, pp. 3830–3834, ISCA.
- [13] Jonathan Shen, Ruoming Pang, Ron J Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, Rj Skerrv-Ryan, et al., “Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions,” in *Proc. ICASSP*, 2018, pp. 4779–4783.
- [14] Jan K Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio, “Attention-based models for speech recognition,” *NeurIPS*, pp. 577–585, 2015.
- [15] Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae, “Hifi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis,” *NeurIPS*, vol. 33, pp. 17022–17033, 2020.
- [16] Triantafyllos Afouras, Joon Son Chung, and Andrew Senior, “LRS3-TED: a large-scale dataset for visual speech recognition,” *arXiv preprint arXiv:1809.00496*, 2018.
- [17] Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N Chang, Sungbok Lee, and Shrikanth S Narayanan, “IEMO-CAP: Interactive emotional dyadic motion capture database,” *Language Resources and Evaluation*, vol. 42, no. 4, pp. 335–359, 2008.
- [18] Lili Guo, Longbiao Wang, Chenglin Xu, Jianwu Dang, Eng Siong Chng, and Haizhou Li, “Representation learning with spectro-temporal-channel attention for speech emotion recognition,” in *Proc. ICASSP*, 2021, pp. 6304–6308.
- [19] Lorenzo Tarantino, Philip N Garner, Alexandros Lazaridis, et al., “Self-attention for speech emotion recognition,” in *Proc. INTERSPEECH*. 2019, pp. 2578–2582, ISCA.
- [20] Heqing Zou, Yuke Si, Chen Chen, Deepu Rajan, and Eng Siong Chng, “Speech emotion recognition with co-attention based multi-level acoustic information,” in *Proc. ICASSP*, 2022, pp. 7367–7371.
- [21] Kun Zhou, Berrak Sisman, and Haizhou Li, “Transforming spectrum and prosody for emotional voice conversion with non-parallel training data,” in *Proc. of The Speaker and Language Recognition Workshop*, 2020, p. 230–237.
- [22] Raghavendra Pappagari, Jesús Villalba, Piotr Żelasko, Laureano Moro-Velazquez, and Najim Dehak, “CopyPaste: An augmentation method for speech emotion recognition,” in *Proc. ICASSP*, 2021, pp. 6324–6328.
- [23] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatuo Gu, and Michael Auli, “data2vec: A general framework for self-supervised learning in speech, vision and language,” in *International Conference on Machine Learning*. 2022, vol. 162, pp. 1298–1312, PMLR.
- [24] Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, et al., “WavLM: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, pp. 1505–1518, 2022.
- [25] Caroline Etienne, Guillaume Fidanza, Andrei Petrovskii, Laurence Devillers, and Benoit Schmauch, “CNN+LSTM architecture for speech emotion recognition with data augmentation,” in *Workshop on Speech, Music and Mind (SMM)*, 2018, pp. 21–25.