



Emotion Recognition from Speech Using Wav2vec 2.0 Embeddings

Leonardo Pepino^{1,2}, Pablo Riera¹, Luciana Ferrer¹

¹Instituto de Investigación en Ciencias de la Computación (ICC), CONICET-UBA, Argentina

²Departamento de Computación, FCEyN, Universidad de Buenos Aires (UBA), Argentina

lpepino@dc.uba.ar, priera@dc.uba.ar, lferrer@dc.uba.ar

Abstract

Emotion recognition datasets are relatively small, making the use of deep learning techniques challenging. In this work, we propose a transfer learning method for speech emotion recognition (SER) where features extracted from pre-trained wav2vec 2.0 models are used as input to shallow neural networks to recognize emotions from speech. We propose a way to combine the output of several layers from the pre-trained model, producing richer speech representations than the model's output alone. We evaluate the proposed approaches on two standard emotion databases, IEMOCAP and RAVDESS, and compare different feature extraction techniques using two wav2vec 2.0 models: a generic one, and one finetuned for speech recognition. We also experiment with different shallow architectures for our speech emotion recognition model, and report baseline results using traditional features. Finally, we show that our best performing models have better average recall than previous approaches that use deep neural networks trained on spectrograms and waveforms or shallow neural networks trained on features extracted from wav2vec 1.0.

Index Terms: speech emotion recognition, transfer learning, wav2vec 2.0

1. Introduction

In recent years, progressively more people are interacting with virtual voice assistants, such as Siri, Alexa, Cortana and Google Assistant. Interfaces that ignore a user's emotional state or fail to manifest the appropriate emotion can dramatically impede performance and risks being perceived as cold, socially inept, untrustworthy, and incompetent [1]. Because of this, speech emotion recognition is becoming an increasingly relevant task.

Several datasets have been created over the years for training and evaluating emotion recognition models, including SAVEE [2], RAVDESS [3], EMODB [4], IEMOCAP [5], and MSP-Podcast [6]. With the exception of MSP-Podcast, these datasets are relatively small in size, usually including only a few dozen speakers. In terms of modeling techniques, many different traditional approaches have been proposed, using hidden Markov models (HMMs) [7, 8], support vector machines (SVMs) [9, 10], decision trees [11, 12], and, more recently, deep neural networks (DNNs) [13, 14, 15]. Yet, while DNNs have shown large gains over traditional approaches on tasks like automatic speech recognition (ASR) [16] and speaker identification [17], the gains observed on emotion recognition are limited, likely due to the small size of the datasets.

A common strategy when dealing with small training datasets is to apply transfer learning techniques. One approach for transfer learning is to use a model learned for a certain auxiliary task for which large datasets are available for training to improve robustness for the task of interest for which data is scarce. The model learned on the auxiliary task can be used as feature extractor or fine-tuned to the task of interest, after replacing

some of its final layers. Recently, transfer learning approaches have been explored in the field of speech emotion recognition. In [18], a deep neural network based on transformers [19] is pretrained on LibriSpeech [20] using multiple self-supervised objectives at different time scales. Then, the model is used as a feature extractor or fine-tuned for speech emotion recognition, among other downstream tasks. Similarly, in [21] and [22], a deep encoder is pretrained in a contrastive predictive coding task, and the resulting embeddings are tested in speech emotion datasets. Recently, several models for automatic speech recognition (ASR) which use self-supervised pretraining have been released, including wav2vec [23] and VQ-wav2vec [24]. A few recent studies [25, 26, 27] have successfully applied representations from these models as features for emotion recognition.

In line with those works, in this paper, we explore the use of the wav2vec 2.0 model [28], an improved version of the original wav2vec model, as a feature extractor for speech emotion recognition. The main contributions of our paper are (1) the use of wav2vec 2.0 representations for speech emotion recognition which, to our knowledge, had never been used for this task, (2) a novel approach to leverage information from multiple layers of the wav2vec 2.0 model, which leads to significant improvements over previous approaches, and (3) an analysis of the importance of the different layers in wav2vec 2.0 for the emotion recognition task. Our results are superior to others in the literature for models based only on acoustic information for IEMOCAP and RAVDESS. The code to replicate the results of this paper has been released at <https://github.com/habla-liaa/ser-with-w2v2>.

2. Methods

In our study, we extracted features from two released wav2vec 2.0 models and used them for speech emotion recognition. In this section, we describe the wav2vec 2.0 model, the datasets used for training and evaluation, and the downstream models.

2.1. Wav2vec 2.0

Wav2vec 2.0 [28] is a framework for self-supervised learning of representations from raw audio. The model consists of a **local encoder**, which contains several convolutional blocks and encodes the raw audio into a sequence of embeddings with a stride of 20 ms and a receptive field of 25 ms. Two models have been released for public use, a large one and a base one where the embeddings are 1024- and 768-dimensional, respectively. A **quantization module**, consisting of a Gumbel-Softmax layer [29], is used to quantize the local encoder representation generating the targets for the model. Then, a **contextualized encoder** takes the local encoder representations as input. Its architecture consists of several transformer encoder blocks [19]. The base model uses 12 transformer blocks with 8 attention heads each, while the large model uses 24 transformer blocks with 16 attention heads each.

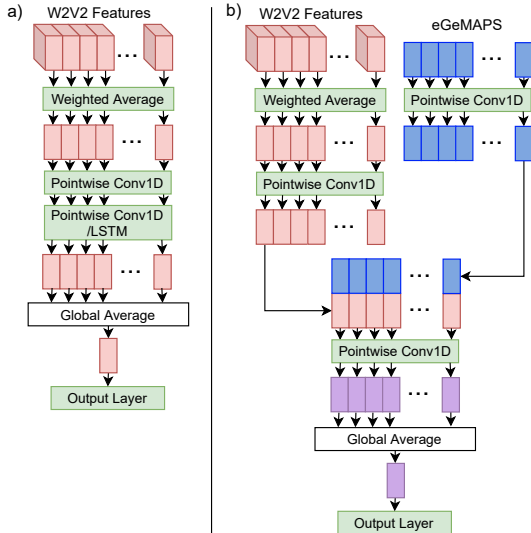


Figure 1: *Downstream models used for this work. The green blocks represent trainable layers. We call the model in figure a) Dense, when using a pointwise convolutional layer, and LSTM, when using an LSTM layer. We call the model in b) the Fusion model. Blocks for each time-step have a height dimension for the features, and an optional depth dimension for the layers in the wav2vec 2.0 model.*

The wav2vec 2.0 model is pretrained in a self-supervised setting, similar to the masked language modelling used in BERT [30] for NLP. Contiguous time steps from the local encoder representations are randomly masked and the model is trained to minimize a contrastive loss between the contextualized encoder outputs and the quantized local encoder representations for those masked time-steps.

The main goal of the wav2vec 2.0 paper was to use the learned representations to improve ASR performance, requiring less data for training and enabling its use for low resource languages. To this end, the model trained as described above is finetuned for ASR using a labelled speech corpus like LibriSpeech. A randomly initialized linear projection is added at the output of the contextual encoder and the connectionist temporal classification (CTC) loss [31] is minimized. The models finetuned in 960 hours of LibriSpeech reach state of the art results in automatic speech recognition when evaluated in different subsets of LibriSpeech. Even when finetuning using considerably less hours, wav2vec 2.0 models reach a performance comparable to the state of the art.

We hypothesize that the information about speech signals learned by the wav2vec 2.0 model might not only be useful for ASR but also for emotion recognition since that same information could allow the model to learn emotion-dependent word or sub-word patterns. Further, these features, particularly the ones from the pre-trained model not fine-tuned for ASR, might also contain prosodic information useful for emotion recognition. In this paper, we compare the performance in speech emotion recognition when using both the wav2vec 2.0 base model pretrained in Librispeech without finetuning¹ (we will call this model Wav2vec2-PT), and a model finetuned for ASR using a 960-hour subset of Librispeech² (Wav2vec2-FT). In both cases, we use the base model as we did not see significant performance

¹https://dl.fbaipublicfiles.com/fairseq/wav2vec/wav2vec_small.pt

²https://dl.fbaipublicfiles.com/fairseq/wav2vec/wav2vec_small.960h.pt

improvements when using the large one, and it allowed us to reduce computational requirements.

2.2. Features

We compare results when using the output of the local and the contextualized encoders as input to the SER models. We also propose to use a weighted average of the outputs of the 12 transformer blocks, along with those of the local encoder. The weights of this average are trained along with the rest of the SER model, as explained in the next subsection.

As baseline features, we use magnitude spectrograms with a hanning window of 25 ms and a hop size of 10 ms, and eGeMAPS [32] low level descriptors (LLD), which are commonly used in the emotion recognition literature. The eGeMAPS features were extracted using opensmile-python [33]. In order to match the sequence lengths of the eGeMAPS features, which use a stride of 10 ms, with the lengths of the wav2vec 2.0 features, which have a stride of 20 ms, we downsampled the eGeMAPS LLDs by averaging every 2 consecutive frames.

All features were normalized by subtracting the mean and dividing by the standard deviation of that feature over all the data for the corresponding speaker. When disabling speaker normalization for comparison of results, we replaced it with global normalization per feature. In this case, the statistics are computed over the training data only, excluding the test data.

2.3. Downstream models

Our downstream model is inspired by [27] due to its simplicity, which should reduce the chance of overfitting. The model we will refer to as *Dense*, consists of 2 dense layers with 128 neurons applied independently to each frame of the input features sequence f , ReLU activation and dropout with a probability of 0.2. This is equivalent to using 1D pointwise convolutional layers (with a kernel size of 1) and 128 filters. The outputs are averaged over time resulting in a vector of size 128. Finally, another dense layer with softmax activation returns probabilities for each of the emotion classes. During the training of the downstream model, the weights of the wav2vec 2.0 model remain unaltered, so it serves as a feature extractor.

For computational reasons, we used a maximum sequence length of 400 for IEMOCAP and 250 for RAVDESS, as input to our SER model. This is equivalent to 8 seconds and 5 seconds, respectively. Note that the output of the contextualized encoder can contain information from the full input waveform which might be longer than the sequence seen by our models. This is because each output of the contextual encoder has a receptive field given by the whole input waveform, since this encoder is a transformer. On the other hand, the local encoder has a limited receptive field, so its outputs can only capture local information.

For the case in which we take as features the activations from both the wav2vec 2.0 local encoder and the transformer blocks, we incorporate a trainable weighted average as the first layer. This layer learns the weights α_i for the vector of activations, f_i , from layer i in the wav2vec 2.0 model, where i goes from 0 to $N = 12$. The vector f_0 corresponds to the local encoder output, f_1 through f_{11} are the outputs of the internal blocks in the transformer and f_{12} is the output of the contextualized encoder. Then, the activations are combined as follows

$$f = \frac{\sum_{i=0}^N \alpha_i f_i}{\sum_{i=0}^N \alpha_i} \quad (1)$$

The weights α_i are initialized with 1.0. Note that the layer ac-

tivations can be combined in this way because they are all the same size. Another possible approach would be to concatenate the outputs from all layers, but the high dimensionality of the resulting feature vectors would be a challenge for the downstream model. For this reason, we decided to use the weighted average, which was inspired by the approach used in [34] for NLP.

In the results section, we compare performance of the Dense model with an LSTM model, in which the second dense layer is replaced by an LSTM layer to account for long-range temporal dependencies. Finally, we also evaluate a third model, called Fusion model, which incorporates a branch taking as input eGeMAPS features. In this last model, the outputs of the first dense layer are concatenated before applying the second dense layer. Our hypothesis is that some prosodic information could be absent from the wav2vec 2.0 representations, making the use of prosodic attributes such as eGeMAPS beneficial. The described downstream models are illustrated in Figure 1.

The downstream models were trained using batches of 32 utterances, Adam optimizer with a learning rate of 0.001, and early stopping with a patience of 4 epochs.

2.4. Datasets

The Ryerson Audio-Visual Database of Emotional Speech and Song (RAVDESS) [3] is a multi-modal database of emotional speech and song. It features 24 different actors (12 males and 12 females) enacting 2 statements: “Kids are talking by the door” and “Dogs are sitting by the door.” with 8 different emotions: happy, sad, angry, fearful, surprise, disgust, calm, and neutral. These emotions are expressed in two different intensities: normal and strong, except for neutral (normal only). Each of the combinations was spoken and sung, and repeated 2 times, leading to 104 unique vocalizations per actor. Following [35], we merged the neutral and calm emotions, resulting in 7 emotions, and used the first 20 actors for training, actors 20-22 for validation to do early stopping, and actors 22-24 for test.

The Interactive Emotional Dyadic Motion Capture (IEMOCAP) dataset [5] has a length of approximately 12 hours and consists of scripted and improvised dialogues by 10 speakers. It is composed of 5 sessions, each including speech from an actor and an actress. Annotators were asked to label each sample choosing one or more labels from a pool of emotions. In this work, we used 4 emotional classes: anger, happiness, sadness and neutral, and following the work in [36], we relabeled excitement samples as happiness. Instances from other classes and with no majority label across the annotations were discarded.³ We also trimmed the waveforms to 15 seconds to reduce the computational requirements when extracting wav2vec 2.0 features. Only 2% of the waveforms were longer than 15 seconds. To evaluate the models in IEMOCAP, we performed 5-fold cross-validation, leaving one session out for each of the folds, as it is a standard practice with this dataset. We used one of the 4 training sessions to perform early stopping for each cross-validation model.

3. Results and discussion

Results are obtained by running the system with 5 different random seeds. Table 1 shows the mean and standard deviation over those seeds of the average recall over all emotion classes ob-

³Note that discarding no-agreement samples and samples from non-target emotions is not an ideal practice [37]. Here, we decided to do this since it is standard practice in emotion recognition literature, facilitating comparisons across papers.

Table 1: Mean and standard deviation over 5 seeds of the average recall (%) over emotions for the different features and models on IEMOCAP and RAVDESS. All results are obtained with the Dense model.

Pretrained model	Features	Dataset	
		IEMOCAP	RAVDESS
None	eGeMAPS	52.4 ± 0.1	57.0 ± 2.4
	Spectrogram	49.8 ± 1.0	44.5 ± 0.8
Wav2vec2-PT	Local enc.	60.3 ± 0.7	65.4 ± 1.7
	Cont. enc.	58.5 ± 0.6	69.0 ± 0.2
	All layers	67.2 ± 0.7	84.3 ± 1.7
Wav2vec2-FT	Local enc.	57.3 ± 1.0	58.8 ± 2.7
	Cont. enc.	44.6 ± 1.0	37.5 ± 3.0
	All layers	63.8 ± 0.3	68.7 ± 0.9

tained using the different features extracted from wav2vec 2.0 models. We also show two systems based on eGeMAPS and spectrogram features, which can be considered as baselines. In all cases, features are normalized by speaker and the Dense model architecture in Figure 1 is used as downstream model. We can see that for both wav2vec 2.0 models, the one finetuned for ASR (wav2vec2-FT) and the one that is not finetuned (wav2vec2-PT), the local encoder representations lead to better results than both of the baseline features. It is worth noting that eGeMAPS, spectrograms and the wav2vec 2.0 local encoder representations contain information restricted only to a local window around each frame, of 60 ms for eGeMAPS, and 25 ms for the others. Also, the downstream model combines information from consecutive frames using just global average, which is a very simple approach that might be suboptimal because it cannot take into account temporal patterns in the features. In spite of that, the local encoder representations, particularly the ones obtained from the PT model, reach a performance comparable to much more complex models like the one proposed in [14], which is a CNN with Bi-LSTM layers trained in a fully supervised setting with spectrograms as input (see Table 2).

We can also see that the wav2vec2-PT features perform better than wav2vec2-FT features in all cases. In particular, the model using only the contextualized encoder outputs for the wav2vec2-FT model has the worst results in the table. We hypothesize that this is because when the model is finetuned for an ASR task, information that is not relevant for that task but might be relevant for speech emotion recognition is lost from the embeddings. For example, information about the pitch might not be important for speech recognition, while it is essential for speech emotion recognition.

Further, Table 1 shows that using a weighted average of the transformer blocks outputs along with the local encoder outputs (rows labelled “All layers”) results in better performance than using only the local or the contextual encoder outputs. Using only the local encoder representations might not give information about events occurring at time scales larger than the receptive field (25 ms). Further, the output of the contextual encoder might be similar to the quantized local encoder representations, since the contrastive loss objective is designed to achieve this goal. Hence, using both of these layers, along with all intermediate layers in the transformer provides additional information that is valuable to the model.

Figure 2 shows the weights α_i from the weighted average layer once the downstream models are trained. We can see that the middle layers are given larger weights. This could be because these layers have already contextualized enough informa-




Figure 2: Weights of the trainable weighted average layer of the Dense model. Index 0 corresponds to the local encoder output. Indices 1 through 12 correspond to the outputs of the layers in the contextual encoder, with layer 12 being the output. Solid lines correspond to the wav2vec2-PT features, while dashed lines correspond to the wav2vec2-FT features.

tion from the local encoder, but they are not yet too specific to the pretraining or finetuning task as the last layers. Also, in the case of the feature extracted from the wav2vec2-PT model, the weights tend to be larger in the layers closer to the output when compared with the weights for the features from the wav2vec2-FT model. This again suggests that the layers of the wav2vec2-FT model closer to the output are less useful for emotion recognition than those of the wav2vec2-PT model. Finally, note that in both datasets, the different training seeds lead to very similar weights, as observed from the error bars in Figure 2, from which we can conclude that these weights are not too sensitive to the neural network initialization.

Finally, we experimented with several variations of the best performing model in Table 1. The first line in Table 2 shows the results for that system, corresponding to the fifth line in Table 1. The rest of the lines in Table 2 show results for that model and the others in Figure 1 applying global normalization instead of speaker-dependent normalization. This is the scenario that is most commonly used in papers as is the one with less assumptions about the available data, treating all samples from a speaker as independent from each other, not assuming that additional samples from a speaker are available at test time.

Comparing the first and the second line in Table 2, we can see that results significantly degrade when doing global normalization. This is expected since normalization by speaker helps the model focus on the emotion characteristics by eliminating part of the speaker information. The degradation is larger in RAVDESS probably because, in this dataset, audios from different emotions do not have much variation in lexical content. Hence, by reducing the effect of speaker in the features all that is left is the variation due to the emotions. On the other hand, in IEMOCAP data, the variation due to lexical content is still in the samples after speaker normalization.

The third and fourth lines in Table 2 show the results obtained using the LSTM model, and the Fusion model. The latter fuses eGeMAPS with the wav2vec 2.0 features. We can see that using an LSTM layer before the global pooling, does not seem to bring improvements over using a simple dense layer. This might be because wav2vec 2.0 features are already contextualized and have global information about the full utterance. Also, LSTMs might be more prone to overfitting or optimization problems than a simpler dense layer. Finally, the table shows some modest improvements from the addition of eGeMAPS features, suggesting that wav2vec 2.0 features may be lacking some of the information present in eGeMAPS.

Table 2 also shows some of the results obtained in other

Table 2: Mean and standard deviation over 5 seeds for the average recall over emotions for the model variants in Figure 1 using the wav2vec2-PT features from all layers, with two types of normalization (top block) compared with results from other works, which all use global normalization (bottom block). In bold are the best results using global normalization.

Model - Norm	IEMOCAP	RAVDESS
Dense - Speaker	67.2 \pm 0.7	84.3 \pm 1.7
Dense - Global	65.8 \pm 0.3	75.7 \pm 2.3
LSTM - Global	64.8 \pm 1.9	74.6 \pm 3.7
Fusion - Global	66.3 \pm 0.7	77.5 \pm 1.0
BiLSTM w. attention [13]	58.8	-
CNN-BiLSTM [14]	59.4	-
TDNN-LSTM w. attention [15]	60.7	-
Wav2Vec [27]	64.3	-

works on the IEMOCAP dataset, using the same experimental setup we are using. For a fair comparison with our work, we restrict the comparison to models in the literature that do not use automatic or manual transcriptions. We can see that all of our models perform better than the state of the art.

Finally, we also compare our results with those in [38], which consists of a deep convolutional neural network using acoustic features as input. For this comparison, we used the dense downstream model with global normalization and imitated their experimental setup both for IEMOCAP and RAVDESS. For IEMOCAP, they only use the improvised sessions and full agreement utterances. For RAVDESS, they use 5 fold cross-validation, dividing the data randomly. Moreover, they do not merge calm and neutral emotions, so the total number of emotions to be predicted is 8. We outperformed the models in [38] for both datasets obtaining an average recall of 84.1 \pm 1.2 % in RAVDESS, and 72.1 \pm 0.9% in IEMOCAP, compared to the results in the paper which are 71.6% and 64.3%, respectively.

4. Conclusions

In this work, we explored different ways of extracting and modeling features from pretrained wav2vec 2.0 models for speech emotion recognition. We proposed to combine the different layers in the wav2vec 2.0 model using trainable weights and model the resulting features with a simple DNN with a time-wise pooling layer. We evaluated our models on two standard emotion datasets, IEMOCAP and RAVDESS, and showed superior results on both cases, compared to those in recent literature. We found that the combination of information from different layers in the wav2vec 2.0 model led to improved results over using only the encoder outputs, as in previous works. Further, we found that the combination of the wav2vec 2.0 features with a set of prosodic features gave additional gains, suggesting that the wav2vec 2.0 model does not contain all the prosodic information needed for emotion recognition. Finally, we showed that a wav2vec 2.0 model finetuned for the task of ASR worked worse than the one trained only with the self-supervised task, indicating that the finetuning eliminates information from the embeddings that is useful for emotion recognition.

5. Acknowledgements

This material is based upon work supported by a Google Faculty Research Award, 2019, and an Amazon Research Award, 2019.

6. References

- [1] S. Brave and C. Nass, "Emotion in human-computer interaction," in *Human-computer interaction fundamentals*. CRC Press Boca Raton, FL, USA, 2009, vol. 20094635.
- [2] S. Haq and P. Jackson, "Machine Audition: Principles, Algorithms and Systems," W. Wang, Ed. IGI Global, 2010.
- [3] 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, 2018.
- [4] F. Burkhardt, A. Paeschke, M. Rolfes, W. F. Sendlmeier, and B. Weiss, "A database of german emotional speech," in *Ninth European Conference on Speech Communication and Technology*, 2005.
- [5] 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, 2008.
- [6] R. Lotfian and C. Busso, "Building naturalistic emotionally balanced speech corpus by retrieving emotional speech from existing podcast recordings," *IEEE Transactions on Affective Computing*, vol. 10, no. 4, 2019.
- [7] B. Schuller, G. Rigoll, and M. Lang, *Hidden Markov Model-based Speech Emotion Recognition*, 2003, vol. 2.
- [8] N. Sato and Y. Obuchi, "Emotion recognition using mel-frequency cepstral coefficients," *Information and Media Technologies*, vol. 2, no. 3, 2007.
- [9] P. Shen, Z. Changjun, and X. Chen, "Automatic speech emotion recognition using support vector machine," in *Proceedings of 2011 International Conference on Electronic & Mechanical Engineering and Information Technology*, vol. 2. IEEE, 2011.
- [10] K. S. Rao, S. G. Koolagudi, and R. R. Vempada, "Emotion recognition from speech using global and local prosodic features," *International journal of speech technology*, vol. 16, no. 2, 2013.
- [11] M. Borchert and A. Dusterhoft, "Emotions in speech-experiments with prosody and quality features in speech for use in categorical and dimensional emotion recognition environments," in *International Conference on Natural Language Processing and Knowledge Engineering*. IEEE, 2005.
- [12] C.-C. Lee, E. Mower, C. Busso, S. Lee, and S. Narayanan, "Emotion recognition using a hierarchical binary decision tree approach," *Speech Communication*, vol. 53, no. 9-10, 2011.
- [13] S. Mirsamadi, E. Barsoum, and C. Zhang, "Automatic speech emotion recognition using recurrent neural networks with local attention," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017.
- [14] A. Satt, S. Rozenberg, and R. Hoory, "Efficient emotion recognition from speech using deep learning on spectrograms," in *Interspeech*, 2017.
- [15] M. Sarma, P. Ghahremani, D. Povey, N. K. Goel, K. K. Sarma, and N. Dehak, "Emotion identification from raw speech signals using dnns," in *Interspeech*, 2018.
- [16] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu *et al.*, "Conformer: Convolution-augmented transformer for speech recognition," *arXiv preprint arXiv:2005.08100*, 2020.
- [17] P. Matějka, O. Glembek, O. Novotný, O. Plchot, F. Grézil, L. Burget, and J. H. Cernocký, "Analysis of dnn approaches to speaker identification," in *ICASSP*. IEEE, 2016.
- [18] Y. Zhao, D. Yin, C. Luo, Z. Zhao, C. Tang, W. Zeng, and Z.-J. Zha, "General-Purpose Speech Representation Learning through a Self-Supervised Multi-Granularity Framework," *arXiv:2102.01930*, 2021.
- [19] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *arXiv preprint arXiv:1706.03762*, 2017.
- [20] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An asr corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015.
- [21] M. Li, B. Yang, J. Levy, A. Stolcke, V. Rozgic, S. Matsoukas, C. Papayiannis, D. Bone, and C. Wang, "Contrastive Unsupervised Learning for Speech Emotion Recognition," *arXiv:2102.06357*, 2021.
- [22] R. Zhang, H. Wu, W. Li, D. Jiang, W. Zou, and X. Li, "Transformer based unsupervised pre-training for acoustic representation learning," *arXiv:2007.14602*, 2021.
- [23] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised Pre-training for Speech Recognition," *arXiv:1904.05862*, 2019.
- [24] A. Baevski, S. Schneider, and M. Auli, "vq-wav2vec: Self-supervised learning of discrete speech representations," in *International Conference on Learning Representations*, 2020.
- [25] S. Siriwardhana, A. Reis, R. Weerasekera, and S. Nanayakkara, "Jointly Fine-Tuning "BERT-like" Self Supervised Models to Improve Multimodal Speech Emotion Recognition," *arXiv:2008.06682*, 2020.
- [26] M. Macary, M. Tahon, Y. Estève, and A. Rousseau, "On the use of Self-supervised Pre-trained Acoustic and Linguistic Features for Continuous Speech Emotion Recognition," *arXiv:2011.09212*, 2020.
- [27] J. Boigne, B. Liyanage, and T. Östrem, "Recognizing More Emotions with Less Data Using Self-supervised Transfer Learning," *arXiv:2011.05585*, 2020.
- [28] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations," *arXiv:2006.11477*, 2020.
- [29] E. Jang, S. Gu, and B. Poole, "Categorical reparameterization with gumbel-softmax," *arXiv preprint arXiv:1611.01144*, 2016.
- [30] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [31] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *ICML*, 2006.
- [32] F. Eyben, K. R. Scherer, B. W. Schuller, J. Sundberg, E. André, C. Busso, L. Y. Devillers, J. Epps, P. Laukka, S. S. Narayanan *et al.*, "The geneva minimalistic acoustic parameter set (gemaps) for voice research and affective computing," *IEEE transactions on affective computing*, vol. 7, no. 2, 2015.
- [33] F. Eyben, M. Wöllmer, and B. Schuller, "Opensmile: the munich versatile and fast open-source audio feature extractor," in *Proceedings of the 18th ACM international conference on Multimedia*, 2010.
- [34] M. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, "Deep contextualized word representations," in *NAACL*, 2018.
- [35] K. Venkataramanan and H. R. Rajamohan, "Emotion recognition from speech," *arXiv preprint arXiv:1912.10458*, 2019.
- [36] H. M. Fayek, M. Lech, and L. Cavedon, "Evaluating deep learning architectures for Speech Emotion Recognition," *Neural Networks*, vol. 92, 2017.
- [37] P. Riera, L. Ferrer, A. Gravano, and L. Gauder, "No sample left behind: Towards a comprehensive evaluation of speech emotion recognition system," in *Proc. Workshop on Speech, Music and Mind 2019*, 2019.
- [38] D. Issa, M. Fatih Demirci, and A. Yazici, "Speech emotion recognition with deep convolutional neural networks," *Biomedical Signal Processing and Control*, vol. 59, 2020.