



# Direct Modelling of Speech Emotion from Raw Speech

Siddique Latif<sup>1,2</sup>, Rajib Rana<sup>1</sup>, Sara Khalifa<sup>2,3</sup>, Raja Jurdak<sup>2</sup>, Julien Epps<sup>3</sup>

<sup>1</sup>University of Southern Queensland, Australia

<sup>2</sup>Distributed Sensing Systems Group, Data61, CSIRO Australia

<sup>3</sup>University of New South Wales, Australia

siddique.latif@usq.edu.au

## Abstract

Speech emotion recognition is a challenging task and heavily depends on hand-engineered acoustic features, which are typically crafted to echo human perception of speech signals. However, a filter bank that is designed from perceptual evidence is not always guaranteed to be the best in a statistical modelling framework where the end goal is for example emotion classification. This has fuelled the emerging trend of learning representations from raw speech especially using deep learning neural networks. In particular, a combination of Convolutional Neural Networks (CNNs) and Long Short Term Memory (LSTM) have gained great traction for the intrinsic property of LSTM in learning contextual information crucial for emotion recognition; and CNNs been used for its ability to overcome the scalability problem of regular neural networks. In this paper, we show that there are still opportunities to improve the performance of emotion recognition from the raw speech by exploiting the properties of CNN in modelling contextual information. We propose the use of parallel convolutional layers to harness multiple temporal resolutions in the feature extraction block that is jointly trained with the LSTM based classification network for the emotion recognition task. Our results suggest that the proposed model can reach the performance of CNN trained with hand-engineered features from both IEMO-CAP and MSP-IMPROV datasets.

**Index Terms:** speech emotion, raw speech, convolutional neural networks, Long Short Term Memory .

## 1. Introduction

Automatic speech emotion recognition has many important applications, such as diagnosis of depression [1], distress [2], monitoring mood state for bipolar patients [3], and so on. However, emotion recognition from speech is a complex task as emotional expression could vary significantly due to a multitude of contextual factors including culture, age, gender, accent, surrounding environment and so on [4].

Research on speech emotion recognition primarily focuses on hand-engineered acoustic features as well as on designing efficient machine learning based models for accurate emotion prediction [5]. In particular, building an appropriate feature representation and designing an appropriate classifier for these features have often been treated as separate problems in the speech recognition community. One drawback of this approach is that the designed features might not be best for the classification objective at hand. LogMel features have been the most popular feature to train Deep Neural Networks (DNNs) and their variants to date. The Mel filter bank is inspired by auditory and physiological evidence of how humans perceive speech signals [6]. However, based on the argument in [7] a filter bank that is designed from perceptual evidence is not always guaranteed

to be the best filter bank in a statistical modelling framework where the end goal is emotion classification. These arguments have led to a recent trend in the machine learning community towards deriving a representation of the input signal directly from raw, unprocessed data. The network learns an intermediate representation of the raw input signal automatically that better suits the task at hand and hence lead to improved performance compared to the classical methods.

A challenging issue in emotion recognition from speech is the efficient modelling of long temporal context [8]. This is because emotions are context-dependent [9] and emotion specific information is embedded in the long temporal contexts [10]. LSTM [11] can model a long range of contexts due to the presence of a special structure called the memory cell. This is why researchers frequently use LSTM for speech emotion recognition [12]. Interestingly, convolutional layer filters can also be used to capture contextual information, which has shown great success in natural language processing (NLP) for sentence classification [13]. In particular, multiple width filters can help improve performance because the model can simultaneously learn multiple contextual dependencies. This has also been validated for speech emotion recognition [14]. However, these studies have generally used multiple width filters in a single layer. Recent studies have shown that parallel convolutional layers can extract temporal information at multiple resolutions from the given data, which can improve the performance of the system [15, 16]. In contrast to using multiple filters in a single layer, in this paper we propose the use of parallel convolutional layers with different filter width to capture diverse contextual information from raw speech.

The key contribution of this paper is the proposed network that consists of a multi-temporal CNN stacked on LSTM. The proposed construct of CNN provides an additional layer for capturing contextual information at multiple temporal resolutions and is designed to complement LSTM for modelling long-term contextual information from raw speech.

## 2. Related Work

Many studies have considered the use of DNN models for processing raw waveforms directly, but the majority of these are in the field of automatic speech recognition (ASR). Dimitri et al. [17] used a framework of CNN for ASR and achieved competitive results to standard short-term spectral features. The authors showed that convolution layers act like a data-driven filterbank and can model spectral envelope of raw speech. The authors in [18] showed that CNN can learn more generalised features across different databases from raw speech compared to artificial neural networks (ANNs) and other feature-based approaches. The complementary approach of using CNN and LSTM jointly for raw speech has been evaluated in [7, 19]. The

authors have shown that the use of LSTM with CNN, helps to reduce the word error rate and achieve competitive performance to the standard feature-based approaches. Besides ASR, researchers have highlighted the feature learning power of different DNN models from raw speech for many other tasks including environmental sound recognition [20], speaker identification [21, 22], and automatic tagging [23].

Few studies have attempted to model emotions using raw speech with results not quite matching feature-based methods. For instance, two studies [24, 12] used end-to-end models by combining CNN and LSTM layers for predicting valence-activation on RECOLA database [25] and achieved promising results. Sarma et al. [10] evaluated time-delay neural network (TDNN) based multiple architectures to model long term dependencies of speech emotion and provided promising results on IEMOCAP dataset. In [14] multi-width filter CNN was applied to hand-engineered features (Mel Filterbanks (MFBs)) which provided competitive results to the systems trained on popular emotional feature sets. In contrast to previous studies [10, 12, 24], we propose a parallel configuration of convolutional layers with multiple filter lengths in feature extraction block to harness multiple temporal resolutions and simultaneously extract multiple contextual dependencies. The classification block is jointly optimised with the feature extraction block to achieve the emotion classification objective.

### 3. Model

Our model consists of two parts: a feature extraction block and a classification block, as shown in Figure 1.

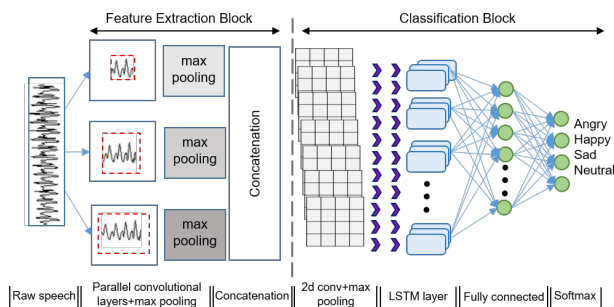


Figure 1: Schematic diagram of proposed model

#### 3.1. Feature extraction block

In the feature extraction block, we use parallel convolutional layers with multiple filter lengths to capture both long-term and short-term interactions directly from raw speech. Given an input utterance, the convolutional layer identifies emotionally salient regions using finite-impulse-response filters, since multiple filters with different lengths can capture diverse contextual dependencies simultaneously from the same region [14].

In our model,  $N$  parallel convolutional layers take  $x_t$  as input of raw speech and create  $N$  different sequences of feature maps by convoluting  $x_t$  with a set of filters of different lengths. The output of convoluting each layer consisting of  $n_w$  filters having widths  $k_w$  and strides  $dw$  are computed using (1).

$$y_i^j = f(b_i + \sum_{k=1}^{k_w} w_k^i x_{dw \times (t-1) + k}) \quad 1 \leq i \leq n_w \quad (1)$$

Here  $f(\cdot)$  is the rectified linear function (ReLU) [26]. Another

important component of our classification block is a nonlinear subsampling layer. For this purpose, we use max pooling layer over time which takes the output of each convolutional layer. Max pooling layer reduces the temporal resolution and selects the most salient features by locally aggregating the feature map of each convolutional layer. We then concatenate the outputs of these three pooling layers to get the features with multiple temporal resolutions and provide that to the classification block.

#### 3.2. Classification Block

We construct our classification block by stacking CNN layer on LSTM. This is motivated by the fact that the performance of LSTM can be improved by feeding it with a good representation [19]. LSTM is specialised to model a long range of contexts due to their gated architecture [27]. Emotion in speech are context-dependent, therefore, the contexts modelling abilities of LSTM are utilised to learn the temporal structure of emotions from the given features maps. We pass the outputs of LSTM to the fully connected layers as it transforms the output of LSTM to a more discriminative space that helps the model for target prediction [19]. In this way, our classification block is jointly empowered by the convolutional layer to capture high-level abstraction, the LSTM layer for long-term temporal modelling, and finally the fully connected layer for learning discriminative representations.

## 4. Experimental Setup

#### 4.1. Dataset

We evaluated our model on two popular datasets: MSP-IMPROV [28] and IEMOCAP [29]. Both of these datasets contain dyadic interactions between actors. We only used audio recordings from these datasets.

##### 4.1.1. IEMOCAP

This corpus contains five sessions, where each session has utterances from two speakers (one male and one female). Overall, there are 10 unique speakers. We used four emotions including angry, happy, neutral and sad. To be consistent with previous studies [14, 30], we merged excitement with happiness and considered one class, happy.

##### 4.1.2. MSP-IMPROV

The MSP-IMPROV dataset contains six sessions, where each session comprises of utterances from two speakers, one male, and one female. There are four emotion categories in MSP-IMPROV: angry, neutral, sad, happy, all were used in the experiments.

#### 4.2. Data Pre-processing and Augmentation

We used the data augmentation to increase the size of the training set. In particular, we created two different copies of each utterance following the approach in [31]. For a given training utterance, we created two versions by applying the speed effect at the factors of 0.9 and 1.1. Sox<sup>1</sup> audio manipulation tool was used for data augmentation. For both datasets, we removed the non-speech intervals at the beginning and end of each utterance as was done in [22].

<sup>1</sup><http://sox.sourceforge.net/>

### 4.3. Model Configuration

We implemented our model using Tensorflow library. In the front end, we selected three parallel striding convolutional layers with different filter widths using the validation data. We used one layer with filter window of 25ms with a shift of 10ms to match the standard frame size of emotional feature extraction process. Smaller and larger filters (than the standard) can also extract useful information from raw speech using CNNs [17, 21]. Therefore, we also used two other layers with filter sizes 15ms and 100ms. The filter sizes were chosen using validation data. We used 40 filters in all three layers. We applied max pooling layer after each convolutional layer to extract the most descriptive features. The feature extraction block was jointly optimised with the classification block where we used a combination of CNN and LSTM. First layer of classification block was a 2d convolutional layer, with filter size (2,2) and filter number 32, followed by the max pooling layer with the pooling size (2,2). The feature maps were then given to the LSTM layer with 128 cells for temporal modelling. Finally, we used one fully connected layer with 1024 units before the softmax layer.

Before applying non-linearity in each convolutional layer, we used batch normalisation (BN) [32] layers to alleviate the problem of exploding and vanishing gradients. For regularisation, we used dropout layer after LSTM layer, with a dropout rate of 0.3. We randomly initialised the weights of our network following the techniques in [33]. Similar to [14] we trained all models using the training set, and validation set was used for hyper-parameter selection. For minimisation of cross-entropy loss function, we used RMSProp optimiser [34], with an initial learning rate of  $10^{-4}$ . If the UAR on the validation set did not improve after 5 epochs, we halved the learning rate. We stopped the process if the UAR did not improve for 20 consecutive epochs. For each model used in this work, we repeated the evaluation 10 times and averaged their predictions.

## 5. Experiments and Results

This section reports the experimental validation of the proposed model for speech emotion recognition. We used leave-one-speaker-out scheme for both datasets and report unweighted average recall (UAR) for both datasets. UAR is a widely used metric for speech emotion recognition due to class imbalanced datasets. In each session, we used utterances from one speaker for testing and utterances from the other speaker for validation and early stopping [14]. The remaining utterances from all speakers were used for training the model. For fair comparison with [14], we used the same data augmentation technique [31] to increase the size of the training set.

For baseline results, we trained SVMs using well-known feature sets, such as, MFCC, LogMel, GeMAPS and eGeMAPS [35] for emotion classification. These features were extracted using openSmile toolkit [36]. We used an RBF kernel and performed grid search using validation data to pick the optimal hyper-parameters. For a fair comparison with our model, we used the same augmented data for all SVM experiments.

Table 1 shows the comparison of results using different methods, and also shows the comparison with previous studies [14, 10, 37]. A direct comparison with some of these studies is not possible due to the difference in data augmentation methods used in the studies, which may affect the results. For example, [10] used different data augmentation scheme, while Gong et al. [37] did not use any data augmentation. We, therefore, compare

Table 1: UAR (%) comparison among different models and proposed approach on raw speech

Method	UAR (%)	
	IEMOCAP	MSP-IMPROV
SVM+MFCC	57.15±2.1	52.38±3.7
SVM+LogMel	58.16±2.6	52.54±3.1
SVM+GeMAPS	57.92±3.2	52.10±3.9
SVM+eGeMAPS	58.76±2.6	52.41±4.6
CNN+MFBs [14]	61.8 ±3.0	52.6 ± 3.8
Proposed+ raw	60.23±3.2	52.43 ±4.1
TDNN-LSTM+ raw (no aug) [10]	48.84	—
SimpleNet-CNN+ raw (no aug) [37]	52.9	—
Proposed+ raw (no aug)	56.72±3.3	48.54 ±3.8

our results with these studies [10, 37] without any data augmentation and separate from other results using a double line in Table 1.

## 6. Analysis and Discussion

This section presents an analysis of the proposed model with architectural changes to gain more insight into the results.

### 6.1. Convolutional Layers Analysis

The convolutional layers play a crucial role in the performance of emotion recognition from raw speech [10, 22]. It is interesting to see the effect of using parallel convolutional layers for capturing multi-temporal resolution features from the raw speech in the feature extraction block. We evaluated different number (1,2,3,4) of parallel convolutional layers and reported the associated results in Table 2. The results show that the pro-

Table 2: Effect of using different number of parallel convolutional layers

Layers	UAR (%)	
	IEMOCAP	MSP-IMPROV
1	57.36±2.3	48.36±3.1
2	58.32±2.8	50.12±3.5
3	<b>60.23±3.2</b>	<b>52.43±4.1</b>
4	59.13±3.1	52.21±4.0

posed multi-temporal resolution model with parallel convolutional layers outperforms the single layer architecture. The best results are obtained using 3 parallel layers, which suggests that a suitable number of parallel layers needs to be determined empirically for specific problems.

### 6.2. Pooling Strategies

Since the pooling layer is used for generalisation of time-domain averaging [7], we evaluated three different pooling operations including max,  $l_2$  and average. Table 3 shows that max pooling outperformed others. All results reported in this paper, therefore, use max pooling.

Table 3: UAR (%) with different pooling strategies.

Pooling	UAR (%)	
	IEMOCAP	MSP-IMPROV
max	<b>60.23±3.2</b>	<b>52.43±4.1</b>
$l_2$	59.72±2.8	50.25±3.0
Average	59.50±3.0	51.94 ±3.2

### 6.3. Analysing Classification Block

In this section, we analyse the effect of using different type of layers in the classification block. Results are reported in Table 4 for both datasets using different configuration of layers in the classification block. In all these setups, we use the same feature extraction block consisting of three parallel convolutional layers that provide multiple temporal dependencies. The architectural changes are only made on the classification block. We trained different configuration of classification blocks including three DNNs (1024-512-512), two LSTMs (256-256), and three CNN layers (256 feature map). We also evaluated other combinations, such as, LSTM-DNN (2 LSTM, 1 DNN), CNN-DNN (2 CNN, 1 DNN), CNN-LSTM (1 CNN, 2 LSTM), and CNN-LSTM-DNN (1 CNN, 1 LSTM, and 1 DNN). All these combinations were trained using the evaluation recipe described in Section 4.3.

Table 4: Analysis of results (UAR) using different combination of layers in classification block

Method	UAR (%)	
	IEMOCAP	MSP-IMPROV
DNN	53.36±2.0	48.36±3.2
LSTM-DNN	56.32±2.6	49.58±3.0
LSTM	58.72±2.9	51.21±3.4
CNN-DNN	58.43±2.8	50.44±3.1
CNN-LSTM	59.23±3.0	52.36±3.6
CNN-LSTM-DNN	<b>60.23±3.2</b>	<b>52.43±4.1</b>
CNN	58.52±2.6	50.84±3.6

In Table 4, we observe that only using DNN layers in the classification block hurts the performance of the model. However, their combination with LSTM and CNN is beneficial for the model performance. We achieved the best performance with the classification block while using the combination of CNN, LSTM and DNN layer. This shows that our proposed construct of the classification block, where convolutional layer captures high-level abstraction, LSTM layer performs long-term temporal modelling, and the fully connected layer performs discriminative representations, offers improvements in the performance of emotion recognition.

### 6.4. Input Length Analysis

We evaluated the performance of our proposed model for different signal lengths. Results are presented in Figure 2. Signal




Figure 2: System performance with different signal length on both datasets.

length is an important aspect since short signals create many small segments which need to be merged, whereas long signals can potentially cause buffer overflow in embedded systems with limited memory. For both IEMOCAP and MSP-IMPROV, UAR

increases with the increase in speech signal length in general. When the signal length is small (1 or 2 seconds) emotion recognition can be performed with a small accuracy loss. However, we observe that speech utterance signal of 6 seconds offers the best UAR for both datasets.

### 6.5. Classification Performance Analysis

We compare the results of our proposed model with that of SVM trained on widely used state-of-art feature sets including MFCC, LogMel, GeMAPS, eGeMAPS in Table 1. It can be observed that our proposed architecture modelling multi-temporal feature from raw speech can achieve better performance compared to the powerful classifier SVM using state-of-the-art features.

We also compare our results with three relevant studies [14, 10, 37] and present the results in Table 1. In [14] authors used Mel Filterbanks (MFBS) features as the input to CNNs and showed that CNNs with these hand-engineered features can produce competitive results to the popular feature sets. In contrast, we used raw speech as input to the model and jointly optimised the feature extraction with the classification network. We achieved comparable results with this study on both datasets as presented in Table 1. This shows that capturing multi-temporal dependencies from the raw speech using parallel CNN layers helps to achieve comparable performance to CNNs trained on hand-engineered features. Two other recent studies [10, 37] used raw speech and evaluated their approach on IEMOCAP dataset. We are achieving better results compared to them when data augmentation is not used. Other recent studies [15, 38] used CNN based models and achieved UAR of 61.9% and 61.7% on IEMOCAP dataset using spectrograms as the input. Compared to these studies, we are achieving 60.23% directly using raw speech.

## 7. Conclusions

In this paper, using two widely used emotion corpus: IEMOCAP and MSP-IMPROV, we show that the proposed network of parallel multi-layer CNN stacked on an LSTM offers, (1) better accuracy when compared to existing methods using raw speech waveform for emotion recognition and (2) comparable accuracy to existing methods using state-of-the-art hand-engineered features. We claim that our proposed construct of CNN having parallel convolutional layers with multiple filter lengths capture both long-term and short-term interactions and help us achieve this performance. In our future studies, we aim to further investigate ways to improve emotion recognition accuracy using raw speech. We also aim to perform run-time and computational complexity comparisons between methods using raw-speech and hand-engineered features and report the accuracy-complexity trade-off.

## 8. References

- [1] Y. Zhu, Y. Shang, Z. Shao, and G. Guo, "Automated depression diagnosis based on deep networks to encode facial appearance and dynamics," *IEEE Transactions on Affective Computing*, vol. 9, no. 4, pp. 578–584, 2018.
- [2] R. Rana, S. Latif, R. Gururajan, A. Gray, G. Mackenzie, G. Humphris, and J. Dunn, "Automated screening for distress: A perspective for the future," *European Journal of Cancer Care*, p. e13033, 2019.
- [3] J. Gideon, E. M. Provost, and M. McInnis, "Mood state prediction from speech of varying acoustic quality for individuals with bipo-

- lar disorder,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 2359–2363.
- [4] R. Rana, M. Hume, J. Reilly, R. Jurdak, and J. Soar, “Opportunistic and context-aware affect sensing on smartphones,” *IEEE Pervasive Computing*, vol. 15, no. 2, pp. 60–69, Apr 2016.
  - [5] B. W. Schuller, “Speech emotion recognition: Two decades in a nutshell, benchmarks, and ongoing trends,” *Communications of the ACM*, vol. 61, no. 5, pp. 90–99, 2018.
  - [6] S. Davis and P. Mermelstein, “Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences,” *IEEE transactions on acoustics, speech, and signal processing*, vol. 28, no. 4, pp. 357–366, 1980.
  - [7] T. N. Sainath, R. J. Weiss, A. Senior, K. W. Wilson, and O. Vinyals, “Learning the speech front-end with raw waveform cldnns,” in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
  - [8] L. Tian, J. Moore, and C. Lai, “Recognizing emotions in spoken dialogue with hierarchically fused acoustic and lexical features,” in *2016 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2016, pp. 565–572.
  - [9] S. Latif, R. Rana, J. Qadir, and J. Epps, “Variational autoencoders for learning latent representations of speech emotion: A preliminary study,” in *Proc. Interspeech 2018*, 2018, pp. 3107–3111.
  - [10] M. Sarma, P. Ghahremani, D. Povey, N. K. Goel, K. K. Sarma, and N. Dehak, “Emotion identification from raw speech signals using dnns,” *Proc. Interspeech 2018*, pp. 3097–3101, 2018.
  - [11] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
  - [12] P. Tzirakis, J. Zhang, and B. W. Schuller, “End-to-end speech emotion recognition using deep neural networks,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5089–5093.
  - [13] Y. Zhang and B. Wallace, “A sensitivity analysis of (and practitioners’ guide to) convolutional neural networks for sentence classification,” *arXiv preprint arXiv:1510.03820*, 2015.
  - [14] Z. Aldeneh and E. M. Provost, “Using regional saliency for speech emotion recognition,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 2741–2745.
  - [15] P. Yenigalla, A. Kumar, S. Tripathi, C. Singh, S. Kar, and J. Vepa, “Speech emotion recognition using spectrogram & phoneme embedding,” *Proc. Interspeech 2018*, pp. 3688–3692, 2018.
  - [16] S. Zhang, S. Zhang, T. Huang, and W. Gao, “Speech emotion recognition using deep convolutional neural network and discriminant temporal pyramid matching,” *IEEE Transactions on Multimedia*, vol. 20, no. 6, pp. 1576–1590, 2018.
  - [17] D. Palaz, M. Magimai-Doss, and R. Collobert, “Analysis of cnn-based speech recognition system using raw speech as input,” in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
  - [18] D. Palaz, M. M. Doss, and R. Collobert, “Convolutional neural networks-based continuous speech recognition using raw speech signal,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 4295–4299.
  - [19] T. N. Sainath, O. Vinyals, A. Senior, and H. Sak, “Convolutional, long short-term memory, fully connected deep neural networks,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 4580–4584.
  - [20] W. Dai, C. Dai, S. Qu, J. Li, and S. Das, “Very deep convolutional neural networks for raw waveforms,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 421–425.
  - [21] H. Muckenhirn, M. M. Doss, and S. Marcell, “Towards directly modeling raw speech signal for speaker verification using cnns,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4884–4888.
  - [22] M. Ravanelli and Y. Bengio, “Speech and speaker recognition from raw waveform with sincnet,” *arXiv preprint arXiv:1812.05920*, 2018.
  - [23] S. Dieleman and B. Schrauwen, “End-to-end learning for music audio,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 6964–6968.
  - [24] G. Trigeorgis, F. Ringeval, R. Brueckner, E. Marchi, M. A. Nicolaou, B. Schuller, and S. Zafeiriou, “Adieu features? end-to-end speech emotion recognition using a deep convolutional recurrent network,” in *2016 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2016, pp. 5200–5204.
  - [25] F. Ringeval, A. Sonderegger, J. Sauer, and D. Lalanne, “Introducing the recola multimodal corpus of remote collaborative and affective interactions,” in *2013 10th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*. IEEE, 2013, pp. 1–8.
  - [26] M. D. Zeiler, M. Ranzato, R. Monga, M. Mao, K. Yang, Q. V. Le, P. Nguyen, A. Senior, V. Vanhoucke, J. Dean *et al.*, “On rectified linear units for speech processing,” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2013, pp. 3517–3521.
  - [27] S. Latif, M. Usman, R. Rana, and J. Qadir, “Phonocardiographic sensing using deep learning for abnormal heartbeat detection,” *IEEE Sensors Journal*, vol. 18, no. 22, pp. 9393–9400, 2018.
  - [28] C. Busso, S. Parthasarathy, A. Burmania, M. AbdelWahab, N. Sadoughi, and E. M. Provost, “Msp-improv: An acted corpus of dyadic interactions to study emotion perception,” *IEEE Transactions on Affective Computing*, vol. 8, no. 1, pp. 67–80, 2017.
  - [29] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, “Iemocap: Interactive emotional dyadic motion capture database,” *Language resources and evaluation*, vol. 42, no. 4, p. 335, 2008.
  - [30] S. Latif, R. Rana, S. Younis, J. Qadir, and J. Epps, “Transfer learning for improving speech emotion classification accuracy,” in *Proc. Interspeech 2018*, 2018, pp. 257–261.
  - [31] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, “Audio augmentation for speech recognition,” in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
  - [32] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” *arXiv preprint arXiv:1502.03167*, 2015.
  - [33] K. He, X. Zhang, S. Ren, and J. Sun, “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1026–1034.
  - [34] T. Tieleman and G. Hinton, “Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude,” *COURSERA: Neural networks for machine learning*, vol. 4, no. 2, pp. 26–31, 2012.
  - [35] 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, pp. 190–202, 2016.
  - [36] F. Eyben, F. Wengler, F. Gross, and B. Schuller, “Recent developments in opensmile, the munich open-source multimedia feature extractor,” in *Proceedings of the 21st ACM international conference on Multimedia*. ACM, 2013, pp. 835–838.
  - [37] Y. Gong and C. Poellabauer, “How do deep convolutional neural networks learn from raw audio waveforms?” 2018.
  - [38] C. Etienne, G. Fidanza, A. Petrovskii, L. Devillers, and B. Schmauch, “Cnn+ lstm architecture for speech emotion recognition with data augmentation,” in *Proc. Workshop on Speech, Music and Mind 2018*, 2018, pp. 21–25.