

FRAUG: A FRAME RATE BASED DATA AUGMENTATION METHOD FOR DEPRESSION DETECTION FROM SPEECH SIGNALS

Vijay Ravi ^{1*}, Jinhan Wang ¹, Jonathan Flint ², Abeer Alwan ¹

¹ Dept. of Electrical and Computer Engineering, University of California, Los Angeles, USA
² Dept. of Psychiatry and Biobehavioral Sciences, University of California, Los Angeles, USA

ABSTRACT

In this paper, a data augmentation method is proposed for depression detection from speech signals. Samples for data augmentation were created by changing the frame-width and the frame-shift parameters during the feature extraction process. Unlike other data augmentation methods (such as VTLP, pitch perturbation, or speed perturbation), the proposed method does not explicitly change acoustic parameters but rather the time-frequency resolution of frame-level features. The proposed method was evaluated using two different datasets, models, and input acoustic features. For the DAIC-WOZ (English) dataset when using the DepAudioNet model and mel-Spectrograms as input, the proposed method resulted in an improvement of 5.97% (validation) and 25.13% (test) when compared to the baseline. The improvements for the CONVERGE (Mandarin) dataset when using the x-vector embeddings with CNN as the backend and MFCCs as input features were 9.32% (validation) and 12.99% (test). Baseline systems do not incorporate any data augmentation. Further, the proposed method outperformed commonly used data-augmentation methods such as noise augmentation, VTLP, Speed, and Pitch Perturbation. All improvements were statistically significant.

Index Terms— data augmentation, depression detection, frame rate, time-frequency resolution, x-vector

1. INTRODUCTION

Major depressive disorder (MDD), is a common and serious medical illness that negatively affects how one feels, thinks and acts. At its worst, depression can lead to suicide and death. Globally, more than 264 million people are affected by MDD [1] and by 2030, it is projected to be the second leading cause of disability [2]. However, only a small percent of these cases get diagnosed and an even smaller percent of them gets treated [3]. Automatic systems for MDD assessment can help reduce diagnostic inequality and will allow for early diagnoses. A possible source of information for building such objective screening mechanisms is the human voice. Among others, the speech signal is an important bio-marker of our mental state [4, 5] and can be collected remotely, in a non-invasive manner with no expert supervision.

Recently, speech-based automatic diagnosis of depression has gained significant momentum [6, 7, 8] and advancements in deep learning have pushed their performance to newer heights [9, 10, 11, 12, 13, 14]. However, data scarcity still remains one of the major challenges in building reliable systems for MDD modeling purposes. Given the sensitivity of mental healthcare data, collection of data can be expensive and challenging. Therefore, there is a need to adopt

data augmentation strategies to increase the amount of training data. However, conventional data augmentation techniques (e.g., Vocal-Tract Length Perturbation - VTLP - [15], Speed and Pitch Perturbation [16]) can be counter-productive when applied to para-linguistic applications such as depression detection because they distort the acoustic data and can lose useful information related to the underlying health condition.

Previously, Generative Adversarial Network (GAN) [17] based data augmentation was proposed for depression detection [12]. However, GANs themselves require significant amount of training data to be effective. In [10], spectrograms were rotated to generate new samples and in [14], noise, pitch-shifting and speed perturbation was employed to augment training data. In [13], a multi-window data augmentation was proposed for emotion recognition which used multiple frame-widths. However, the methods proposed in [10, 13, 14] were not compared with conventional data augmentation techniques and were only evaluated using one model.

In contrast, in this paper, a frame rate based data augmentation technique is proposed specifically for the task of depression detection from speech signals. New feature samples were created by varying the frame-width as well as the frame-shift during feature extraction. By changing the frame-rate parameters, the model was provided with different sets of time-frequency resolutions during the training stage. This ensured that acoustic parameters which are thought to correlate with the mental state of the speaker (eg: pitch, formant frequencies, speaking-rate etc. [4, 7]) were not inadvertently modified. Additionally, the proposed method was shown to outperform two commonly used data augmentation methods and was validated on two different datasets, input acoustic features and models.

The remainder of this paper is organized as follows: the proposed multi-frame-rate data augmentation approach is introduced in Section 2. Experimental details of the datasets and models used are described in Section 3. Results are reported and discussed in Section 4 and the conclusion and future directions are provided in Section 5.

2. PROPOSED METHOD

In this section, we describe the proposed data augmentation technique, FraUG. Given an input speech signal $x[n]$, the windowing and feature extraction process for spectral features, such as spectrograms and mel-frequency cepstral coefficients (MFCC), can be represented as:

$$X_r[k] = \sum_{m=0}^{L-1} x[m]w[rR - m]e^{-j(2\pi k/N)m}, \quad (1)$$

where, $w[rR - m]$ is the sliding window, $r \in \mathbb{Z}$, N is the DFT size,

*Corresponding author: vijaysumaravi@ucla.edu

L is the frame-width and R is the frame-shift [18]. The windows overlap by $O = L - R$. R and O are usually specified as a fraction of L , which is specified in time or number of samples.

Changing the values of L and R and thereby the frame rate, changes the time-frequency resolution of the extracted features. Conventionally, to balance resolutions between time and frequency, the parameters L , R and O are fixed. A Hamming window with $L = 25ms$ and $R = 40%$ (i.e. $R = 10ms$) is commonly used [19].

In FrAUG, given a baseline frame-rate with parameters (L_1, R_1) , we augment the training data with features extracted using multiple frame rates with parameters L_i and R_j where $i, j \in \mathbb{N}$. For example, to perform an 8-fold augmentation, frame-widths of L_2, L_3 and frame-shifts R_2, R_3 are used along with baseline parameters, resulting in 9 different combinations such as $(L_1, R_2), (L_2, R_3), (L_3, R_2)$, etc. Thus, the model is provided with 8 additional time-frequency resolutions in the training stage. The main advantage of the proposed method is that it does not alter vocal tract or voice source parameters, particularly those useful for depression detection, and is independent of the dataset and model used. Since the underlying mechanism of windowing is the same for non-spectral features (such as the prosodic features), FrAUG can be extended to other acoustic features as well.

3. DATASETS AND MODELS

The proposed method was applied on two different models using two distinct datasets and two different input acoustic features. The two models are - DepAudioNet [9] which was trained using mel-spectrograms, and a pretrained x-vector embedding generator trained on MFCCs followed by a convolutional neural network (CNN) backend, for the DAIC-WOZ (English, [20]) and the CONVERGE (Mandarin, [21]) datasets, respectively. The datasets and the corresponding models used in this paper are described in this section.

3.1. Datasets

3.1.1. DAIC-WOZ

The Distress analysis interview corpus wizard of Oz (DAIC-WOZ) [20] database comprises audio-visual interviews of 189 participants, male and female, who underwent evaluation of psychological distress. Each participant was assigned a self-assessed depression score through the patient health questionnaire (PHQ-8) method [22]. Audio data belonging only to the participants were extracted using the time-labels provided with the dataset. Recordings from session numbers 318, 321, 341 and 362 were excluded from the training set because of time-labelling errors. This dataset consists of a total of 58 hours of audio data. The data partitioning between train, validation and test sets is the same as that provided with the database description, which is 60%, 20% and 20%, respectively.

3.1.2. CONVERGE

The second depression database used in this paper is in Mandarin and was developed as a part of the China, Oxford and Virginia Commonwealth University Experimental Research on Genetic Epidemiology (CONVERGE) study [21]. The CONVERGE study focused on subjects with increased genetic risk for MDD, and to obtain a more genetically homogeneous sample, only women participants were recruited. Each participant was interviewed by a trained interviewer. The diagnoses of depressive (dysthymia and MDD) disorders were made with the Composite International Diagnostic

Interview (Chinese version) [23], which classifies diagnoses according to the Diagnostic and Statistical Manual of Mental Disorders fourth edition (DSM-IV) criteria. The database includes recordings of the interviews from 3742 individuals classified as suffering from MDD and 4217 healthy individuals. All audio recordings were collected with a sampling rate of 16kHz. The database was randomly split into 60%, 20%, and 20% for the train, validation, and test sets, respectively. This database contains a total of 391 hours of audio data and is characterized by a large degree of phonetic and content variability.

3.2. Models

3.2.1. DepAudioNet

DepAudioNet is a deep neural network model for detecting depression from speech [9]. This model consisted of a one-dimensional convolutional layer (*Conv1D*) and two unidirectional long short-term memory (LSTM) layers [24]. The *Conv1D* layer had 40-dimensional mel-spectrograms as input and a kernel size of 3 with no time-dilation. The *Conv1D* was followed by ReLU non-linearity, a dropout layer with $p_{dropout} = 0.05$ and a max-pooling layer with kernel size 3. This was followed by two 128-dimensional LSTM layers. Finally, a fully connected layer with Sigmoid activation was applied to generate the binary prediction.

To mitigate class imbalance between depression and non-depression classes, training data were pre-processed by random cropping and random sampling [9]. First, each utterance was randomly cropped to fragments which were equal to the length of the shortest utterance in the DAIC-WOZ database. This was done in order to negate any bias towards longer durations of audio samples. Each randomly cropped fragment was segmented into 120 frames, where each frame varied between 32ms to 128ms depending on the frame rate being applied. A training subset was generated by randomly sampling, without replacement, an equal number of depression and non-depression segments.

Each experiment consisted of training five separate models for 100 epochs each using a randomly generated training subset. The final prediction was obtained by averaging the probabilities predicted by the five models. The pre-processing step was applied irrespective of the type of data augmentation used and hyperparameters such as batch size, learning rate and the learning rate reduction factor were tuned for every experiment individually.

This model was used with the DAIC-WOZ dataset. For the baseline experiments, mel-spectrograms were extracted with frame rate parameters of $L = 64ms$, and $R = 50%$, as proposed in [9]. When FrAUG was applied, training data was augmented with up to 8 folds using additional frame rates with parameters $L = 32ms$ and $L = 128ms$ and an overlap R of 25% and 10%. The augmentation frame rates were chosen empirically. Even when augmentation was applied, mel-spectrograms for test and validation sets were extracted at the baseline frame rate.

3.2.2. x-vector embedding with CNN Classifier

An x-vector embedding extractor [25] was pre-trained using CN-Celeb [26], a Mandarin speaker ID dataset. A Kaldi recipe was followed for training the x-vector model [27]. Inputs to the x-vector model were 30-dimensional MFCCs and the model consisted of 5 layers of time-dilated convolutional networks followed by average pooling and two fully connected layers; the size of the x-vector embedding was 512. No additional pre-processing was applied. The readers are referred to [28] for implementation details.

After pre-training the x-vector model, embeddings for the CONVERGE dataset were generated which were then used to train a downstream network for classifying depression. The downstream network was made up of two CNN layers followed by two fully connected layers. The downstream model was trained for 100 epochs with a fixed learning rate of $1e-4$. Data augmentation was only applied during the training of the downstream network i.e. embeddings were extracted for the augmented depression training data along with the unaugmented validation and test data.

Similar to the previous experiment, x-vector embeddings for the baseline experiments were generated using MFCCs extracted with frame rate parameters of $L = 64ms$ and $R = 50%$. When FrAUG was applied, x-vectors were generated using MFCCs extracted with additional frame rate parameters of $L = 32ms$, $L = 128ms$ and $R = 50%$, $R = 25%$. Augmentation frame rates were chosen empirically and up to 8-fold data augmentation was evaluated. Test and validation set features were always extracted at the baseline frame rates.

4. RESULTS AND DISCUSSION

The effectiveness of the proposed approach is demonstrated in three stages – first, it is shown that training a model with FrAUG is better than the baseline approach of single frame rate training. Then, the performance of the proposed method is compared to conventional data augmentation techniques and lastly, the generalizability of the proposed method is evaluated by applying it to a different dataset with a different backend system and different input acoustic features. Model performance is reported in terms of the F1-score [29] which is the harmonic mean of precision and recall. Statistical significance ($p < 0.05$) was evaluated using the McNemar’s test [30]

4.1. Multi Frame Rate Training

In the first set of experiments, DepAudioNet models were trained on the DAIC-WOZ dataset using different combinations of single frame rate and multiple frame rates. In this work, DepAudioNet was chosen as a baseline mainly because of DepAudioNet’s open-source code [24]. The original paper where DepAudioNet was first proposed [9] did not report test-set results because of the unavailability of ground truth labels and computed the F1-score for predictions at the speaker level and not at the frame level [14] nor the segment level [31]. Unfortunately, there is a lack of consensus regarding the evaluation protocols for the DAIC-WOZ dataset.

Performance comparison of single rate training versus multiple frame rate training on the validation set of the DAIC-WOZ dataset is shown in Table 1. The baseline frame rate of $L = 64ms$, $R = 50%$ has an F1-score of 0.619. This is comparable to the reported F1-scores of 0.610 in prior works [9, 24]. In contrast, the best performing configuration is the one with 5-fold data augmentation with multiple frame rate hyper-parameters of $L = 64ms$, $128ms$ and $R = 50%$, $25%$, $10%$. Higher folds of data augmentation were also evaluated but 5-fold produced the best results.

The best performing system has an F1-score of 0.656, a relative improvement of 5.97% ($p = 4.72e-6$) when compared to the baseline. This best performing configuration is better than any of the single frame rate performances, including when only the frame-widths are manipulated as in [13]. A possible explanation for this result might be that a particular combination of time-frequency resolutions, provided to the model in FrAUG, contains depression-related information that is not available to the model when trained using single frame-rate features.

To evaluate the performance on the test set, the best performing configuration on the validation data was selected and compared with the baseline. From Table 1, the best performing system is the one with $L = 64ms$, $128ms$ and $R = 50%$, $25%$, $10%$. The test-set results comparing the baseline and the best model are shown in Table 2. In this case, the proposed approach has an F1-score of 0.478 compared to the baseline score of 0.382 resulting in an improvement of 25.13% ($p = 5.66e-6$).

The authors of this paper acknowledge that higher baseline F1-scores for the test set have been reported in [31, 14]. However, a comparison cannot be made because, unlike the approach in this paper, those systems either employed a different evaluation protocol such as segment-level predictions or re-partitioned the train-validation-test splits. More importantly, the goal of this paper is to show that FrAUG can provide significant gains over a baseline with no such augmentation.

4.2. FrAUG versus Conventional Data Augmentation Methods

To compare FrAUG with conventional data augmentation techniques, DepAudioNet models were trained using the DAIC-WOZ dataset with FrAUG, noise augmentation [25], VTLP-based augmentation [32, 33], speed perturbation and pitch perturbation [16]. The noise augmentation method was similar to the one used in Kaldi. The MUSAN library [34] was used to augment every utterance with randomly chosen foreground noise samples at a randomly chosen SNR of 0,5,10, or 15 dB [25]. The VTLP augmentation was based on the nlpaug library [33] and the method proposed in [15]. Speed and Pitch perturbation were based on [16] and implemented using the Librosa library. For every augmentation method, up to 8-folds of data augmentation was applied and the best performing configuration, based on the validation set, was selected. The results comparing these augmentation methods are presented in Table 3.

As seen in Table 3, FrAUG is the best performing augmentation strategy for both validation and test sets. On the validation set, FrAUG outperforms noise augmentation by 13.2% ($p = 3.43e-6$), VTLP by 4.1% ($p = 4.92e-6$), speed perturbation by 2.7% ($p = 1.53e-6$) and pitch perturbation by 1.2% ($p = 3.71e-5$). On the test set, the proposed approach is comparable to noise augmentation, is better than VTLP augmentation by 3.7% ($p = 4.95e-6$), speed perturbation and pitch perturbation by 11.1% ($p = 4.88e-6$ and $p = 4.64e-6$, respectively). One possible explanation for these results is that VTLP, speech perturbation and pitch perturbation, alters the spectral shape and therefore might be preserving less information about the depressive state of the speaker. In case of noise augmentation, a domain mis-match between training and validation data (noisy vs clean) may be the reason for degraded performance. This shows that FrAUG can serve as an effective data augmentation strategy for depression detection without interfering with task-related acoustic information.

4.3. Extension to CONVERGE Dataset

To show that the proposed approach is independent of the dataset, the model and the input acoustic feature, it was evaluated on the CONVERGE dataset using embeddings extracted from a pre-trained x-vector system. These embeddings were then used to train a CNN model to classify utterances as cases (depressed) or controls (healthy). 3x, 5x and 8x data augmentation was applied.

The effectiveness of FrAUG as applied to the CONVERGE dataset is evident from the results presented in Table 4. In comparison to the baseline F1-score of 0.676 (validation) and 0.654

Table 1. Results, in terms of F1-score, comparing single frame rate training versus multi frame rate training using DepAudioNet and the DAIC-WOZ Validation set. L and R represent frame-width and frame-shift. * denotes the baseline F1-score. The best F1-score is boldfaced.

$\downarrow L \setminus R \rightarrow$	50%	25%	10%	50%, 25%	50%, 10%	25%, 10%	50%, 25%, 10%
32ms	0.601	0.604	0.569	0.606	0.633	0.562	0.604
64ms	0.619*	0.638	0.587	0.612	0.620	0.599	0.613
128ms	0.648	0.627	0.588	0.618	0.628	0.638	0.616
32ms, 64ms	0.637	0.607	0.579	0.615	0.617	0.623	0.602
64ms, 128ms	0.635	0.612	0.576	0.620	0.625	0.617	0.656
32ms, 128ms	0.623	0.633	0.590	0.647	0.610	0.602	0.615
32ms, 64ms, 128ms	0.626	0.607	0.546	0.655	0.600	0.582	0.596

Table 2. Results, in terms of F1-score, for depression detection on DAIC-WOZ test data comparing baseline performance and the best configuration of FrAUG selected using validation data performance. The best F1-score is boldfaced.

L,R Configuration	Avg F1 Score	Data Augmentation
Baseline	0.382	None
L=64ms, R=50%		
L=64ms, 128ms	0.479	5x
R= 50%, 25%, 10%		

Table 3. Results, in terms of F1-score, for depression detection on the DAIC-WOZ dataset comparing proposed method with conventional data augmentation techniques. The best F1-score is boldfaced.

Augmentation Strategy	Validation	Test	Data Augmentation
Baseline	0.619	0.382	None
Noise [34]	0.579	0.477	7x
VTLP [15]	0.630	0.462	3x
Speed Perturbation [16]	0.639	0.431	3x
Pitch Perturbation [16]	0.648	0.431	5x
FrAUG	0.656	0.479	5x

Table 4. Results, in terms of F1-score, for depression detection on the CONVERGE dataset using x-vector embeddings with a CNN classifier as the backend, with and without FrAUG. The best F1-score is boldfaced.

L,R Configuration	Validation	Test	Data Augmentation
Baseline	0.676	0.654	None
(L=64ms,R=50%)			
L=32ms, 64ms	0.705	0.712	3x
R=50%, 25%			
L=32ms, 64ms	0.720	0.719	5x
R=50%, 25%, 10%			
L=32ms, 64ms, 128ms	0.739	0.739	8x
R=50%, 25%, 10%			

(test), the best performing configuration (8-fold augmentation) has a performance of 0.739 and 0.739, respectively. This is an improve-

ment of 9.32% on the validation set ($p = 5.93e-6$) and 12.99% on the test set ($p = 5.88e-6$). The performance of the model improves consistently with increasing amounts of training data. Even though the downstream model was trained on x-vector embeddings and not on the acoustic features themselves, FrAUG improves the classification performance. This is a rather significant outcome because this shows that FrAUG can be beneficial in improving system performance even when applied to downstream tasks after the pre-training step. An important implication of this result is that FrAUG can be applied irrespective of the model training style - supervised pre-training, training from scratch, etc.

5. CONCLUSION AND FUTURE WORK

In this paper, a data augmentation method, called FrAUG, was proposed for depression detection from speech. Training data were augmented with new feature samples created by varying the frame-width as well as the frame-shift parameters during feature extraction. Thus, the proposed approach did not modify vocal tract or voice source related parameters and hence preserved acoustic information that may be important for MDD modeling purposes. The proposed method of data augmentation performed better than a baseline system with no augmentation and four commonly used data augmentation methods. Lastly, the generalizability of the said method was demonstrated by improvements in classification performance on a different dataset with a different model and different input features.

FrAUG improved the classification performance of DepAudioNet [9] trained using mel-spectrograms on the DAIC-WOZ dataset, and of a downstream network trained with x-vector embeddings generated from a pre-trained model [25] using MFCCs on the CONVERGE dataset. It can therefore be suggested that the proposed method is independent of the dataset, the input acoustic features, the model and the model training style.

Frame rate based data augmentation can be reliably used to increase the amount of training data and might prove to be useful in the development of large-scale MDD screening systems. In the future, FrAUG will be applied to other features such as the Voice Quality features [35], and the fusion of FrAUG with other types of data augmentation techniques will be analyzed. Further, FrAUG will be evaluated for other para-linguistic applications such as emotion recognition, detection of Alzheimer's dementia, etc.

6. ACKNOWLEDGEMENTS

This work was funded by the National Institutes of Health under the award number R01MH122569- Combining Voice and Genetic Information to Detect Heterogeneity in Major Depressive Disorder.

7. REFERENCES

- [1] Spencer L James et al., “Global, regional, and national incidence, prevalence, and years lived with disability for 354 diseases and injuries for 195 countries and territories, 1990–2017: a systematic analysis for the global burden of disease study 2017,” *The Lancet*, vol. 392, no. 10159, pp. 1789–1858, 2018.
- [2] Colin D Mathers and Dejan Loncar, “Projections of global mortality and burden of disease from 2002 to 2030,” *PLoS medicine*, vol. 3, no. 11, pp. e442, 2006.
- [3] Kenneth B Wells, Ron D Hays, et al., “Detection of depressive disorder for patients receiving prepaid or fee-for-service care: results from the medical outcomes study,” *JAMA*, vol. 262, no. 23, pp. 3298–3302, 1989.
- [4] Nicholas Cummins, Stefan Scherer, et al., “A review of depression and suicide risk assessment using speech analysis,” *Speech Communication*, vol. 71, pp. 10–49, 2015.
- [5] Vijay Ravi, Soo Jin Park, et al., “Voice quality and between-frame entropy for sleepiness estimation,” *Interspeech*, 2019.
- [6] Sharifa Alghowinem, Roland Goecke, et al., “Detecting depression: a comparison between spontaneous and read speech,” in *ICASSP*. IEEE, 2013, pp. 7547–7551.
- [7] Amber Afshan, Jinxi Guo, et al., “Effectiveness of voice quality features in detecting depression,” *Interspeech*, 2018.
- [8] S Pavankumar Dubagunta et al., “Learning voice source related information for depression detection,” in *ICASSP*. IEEE, 2019, pp. 6525–6529.
- [9] Xingchen Ma, Hongyu Yang, et al., “Depaudionet: An efficient deep model for audio based depression classification,” in *Proceedings of the 6th international workshop on audio/visual emotion challenge*, 2016, pp. 35–42.
- [10] Lang He and Cui Cao, “Automated depression analysis using convolutional neural networks from speech,” *Journal of biomedical informatics*, vol. 83, pp. 103–111, 2018.
- [11] Amir Harati, Elizabeth Shriberg, et al., “Speech-based depression prediction using encoder-weight-only transfer learning and a large corpus,” in *ICASSP*. IEEE, 2021, pp. 7273–7277.
- [12] Le Yang, Dongmei Jiang, and Hichem Sahli, “Feature augmenting networks for improving depression severity estimation from speech signals,” *IEEE Access*, vol. 8, pp. 24033–24045, 2020.
- [13] Sarala Padi, Dinesh Manocha, and Ram D Sriram, “Multi-window data augmentation approach for speech emotion recognition,” *arXiv preprint arXiv:2010.09895*, 2020.
- [14] Emna Rejaibi, Ali Komaty, et al., “Mfcc-based recurrent neural network for automatic clinical depression recognition and assessment from speech,” *Biomedical Signal Processing and Control*, vol. 71, pp. 103107, 2022.
- [15] 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, p. 21.
- [16] Tom Ko, Vijayaditya Peddinti, Daniel Povey, and Sanjeev Khudanpur, “Audio augmentation for speech recognition,” in *Sixteenth annual conference of the international speech communication association*, 2015.
- [17] Ian Goodfellow, Jean Pouget-Abadie, et al., “Generative adversarial nets,” *Advances in neural information processing systems*, vol. 27, 2014.
- [18] Lawrence Rabiner and Ronald Schafer, *Theory and applications of digital speech processing*, Prentice Hall Press, 2010.
- [19] Daniel Povey, Arnab Ghoshal, et al., “The kaldı speech recognition toolkit,” in *ASRU*. 2011, IEEE Signal Processing Society.
- [20] Michel Valstar, Jonathan Gratch, et al., “Avec 2016: Depression, mood, and emotion recognition workshop and challenge,” in *Proceedings of the 6th international workshop on audio/visual emotion challenge*, 2016, pp. 3–10.
- [21] Yun Li, S Shi, et al., “Patterns of co-morbidity with anxiety disorders in chinese women with recurrent major depression,” *Psychological medicine*, vol. 42, no. 6, pp. 1239–1248, 2012.
- [22] Kurt Kroenke, Tara W Strine, et al., “The phq-8 as a measure of current depression in the general population,” *Journal of affective disorders*, vol. 114, no. 1-3, pp. 163–173, 2009.
- [23] W Ter Smitten, MH and Smeets, RMW and Van den Brink, “Composite international diagnostic interview (CIDI), version 2.1,” *Amsterdam: World Health Organization*, , no. 343, pp. 343–345, 1998.
- [24] Andrew Bailey and Mark D Plumbley, “Raw audio for depression detection can be more robust against gender imbalance than mel-spectrogram features,” *arXiv preprint arXiv:2010.15120*, 2020.
- [25] David Snyder, Daniel Garcia-Romero, et al., “X-vectors: Robust dnn embeddings for speaker recognition,” in *ICASSP*. IEEE, 2018, pp. 5329–5333.
- [26] Yue Fan, JW Kang, et al., “Cn-celeb: a challenging chinese speaker recognition dataset,” in *ICASSP*. IEEE, 2020, pp. 7604–7608.
- [27] Lantian Li, Ruiqi Liu, et al., “Cn-celeb: multi-genre speaker recognition,” 2020.
- [28] Manoj Kumar, Tae Jin-Park, et al., “Designing neural speaker embeddings with meta learning,” 2020.
- [29] N Chinchor, “Muc-4 evaluation metrics in proc. of the fourth message understanding conference 22–29,” 1992.
- [30] Quinn McNemar, “Note on the sampling error of the difference between correlated proportions or percentages,” *Psychometrika*, vol. 12, no. 2, pp. 153–157, 1947, Publisher: Springer.
- [31] Alice Othmani, Daoud Kadoch, et al., “Towards robust deep neural networks for affect and depression recognition from speech,” in *International Conference on Pattern Recognition*. Springer, 2021, pp. 5–19.
- [32] Mingke Xu, Fan Zhang, et al., “Speech emotion recognition with multiscale area attention and data augmentation,” in *ICASSP*. IEEE, 2021, pp. 6319–6323.
- [33] Edward Ma, “Nlp augmentation,” <https://github.com/makcedward/nlpaug>, 2019.
- [34] David Snyder, Guoguo Chen, and Daniel Povey, “Musan: A music, speech, and noise corpus,” *arXiv preprint arXiv:1510.08484*, 2015.
- [35] Soo Jin Park, Caroline Sigouin, et al., “Speaker identity and voice quality: Modeling human responses and automatic speaker recognition,” in *Interspeech*, 2016, pp. 1044–1048.