

# Unsupervised Instance Discriminative Learning for Depression Detection from Speech Signals

Jinhan Wang<sup>1</sup>, Vijay Ravi<sup>1</sup>, Jonathan Flint<sup>2</sup>, Abeer Alwan<sup>1</sup>

<sup>1</sup>Dept. of Electrical and Computer Engineering, University of California, Los Angeles, USA

<sup>2</sup>Dept. of Psychiatry and Biobehavioral Sciences, University of California, Los Angeles, USA

wang7875@g.ucla.edu, vijaysumaravi@ucla.edu, jflint@mednet.ucla.edu, alwan@ee.ucla.edu

## Abstract

Major Depressive Disorder (MDD) is a severe illness that affects millions of people, and it is critical to diagnose this disorder as early as possible. Detecting depression from voice signals can be of great help to physicians and can be done without the need of any invasive procedure. Since relevant labelled data are scarce, we propose a modified Instance Discriminative Learning (IDL) method, an unsupervised pre-training technique, to extract augment-invariant and instance-spread-out embeddings. In terms of learning augment-invariant embeddings, various data augmentation methods for speech are investigated, and time-masking is found to provide the best performance. To learn instance-spread-out embeddings, we explore methods for sampling instances for a training batch (distinct speaker-based and random sampling). It is found that the distinct speaker-based sampling provides better performance than the random one, and we hypothesize that this result is because relevant speaker information is preserved in the embedding. Additionally, we propose a novel sampling strategy, Pseudo Instance-based Sampling (PIS), based on clustering algorithms, to enhance spread-out characteristics of the embeddings. Experiments are conducted with DepAudioNet on DAIC-WOZ (English) and CONVERGE (Mandarin) datasets, and statistically significant improvements are observed in the detection of MDD relative to the baseline with no pre-training.

**Index Terms:** unsupervised pre-training, depression detection, speaker information, instance discriminative learning

## 1. Introduction

Major Depressive Disorder (MDD) is one of the most severe chronic mental health disorders, characterised by persistent depressed mood, loss of interest in social activities, lack of energy, and even suicidal thoughts [1]. Around 4.4% of the world's population is estimated to suffer from depression [2]. It is believed that the cause of MDD is multifactorial, which makes it unexplainable by conventional mechanisms and difficult to diagnose [3]. Conventionally, diagnosing MDD is conducted by questionnaires [4]. However, many subjective factors might cause mis-diagnosis [5]. Therefore, automatic depression detection systems might help alleviate the above mentioned challenges.

There have been several studies focusing on automatic depression detection from signals, such as voice, video, electroencephalogram (EEG), etc. [6, 7, 8, 9]. Among those approaches, automatic depression detection based on speech signals draws more attention because speech signals can be an effective clinical marker for depression [10, 11]. Speech signals also prevail because they are relatively easy to obtain and analyze [6]. Depressed individuals typically have slow, uniform, monotonous and/or hesitant voice characteristics [12, 13]. There have been successful approaches to depression detection based on speech

signals. Various features have been investigated, such as Mel Frequency Cepstral Coefficients (MFCCs) [14, 15], Mel Filterbanks [16], vocal tract coordination (VTC) [17] and voice quality features [18, 19, 10]. Aside from these handcrafted features, some approaches use raw audio signals [20] in an end-to-end manner or high-level features extracted from pre-trained models, like x-vector [21], i-vector [18], NN2Vec [22], and Depression audio embedding (DEPA) [23]. In the back-end domain, researchers apply model architectures based on Convolutional Neural Networks (CNN) and Long-Short-Term Memory (LSTM) to capture spatial and temporal patterns of depression [24, 16]. Other additional mechanisms might be incorporated externally to improve model performance. For example, in [24], an ensemble method with a CNN has shown good robustness to acoustic variability, and in [22] Salekin et al. reformulate the task into a weakly supervised learning problem.

Unsupervised pre-training (UPT) is shown to be effective in various tasks given data scarcity, such as Automatic Speech Recognition (ASR) [25, 26], and Speech Emotion Recognition (SER) [27, 28]. For SER tasks, Contrastive Predictive Coding (CPC) is applied by separating positive examples from negative examples though infoNCE minimization [27]. Speech SimCLR is proposed to optimize contrastive loss between samples that are augmented [28]. There have been a few studies that focused on UPT for depression detection. Most of these approaches use UPT models for feature extraction [22, 23, 29]. However, it's highly likely that the extracted features might lose important information for depression classification, since the pre-trained model is not optimized for depression detection. Hence, there is a domain discrepancy problem which might negatively affect performance. Inspired by a successful UPT method, Instance Discriminative Learning (IDL) [30], for image classification tasks, we propose a modified IDL approach for depression detection based on speech signals to extract augment-invariant and instance-spread-out embeddings in the pre-training stage. We investigate different sampling strategies for pre-training and find that the method that works the best preserves some speaker information. Additionally, a new sampling strategy, Pseudo Instance-based Sampling (PIS), is proposed to boost instance-spread-out characteristics.

The remainder of this paper is organized as follows: Section 2 presents the IDL technique and the proposed PIS method. The experimental setup is described in Section 3, followed by results and discussion in Section 4. Section 5 concludes the paper.

## 2. Instance Discriminative Learning

### 2.1. Pre-training Methodology

The schematic diagram of IDL [30] is shown in Figure 1. The assumption of IDL pre-training is that, the embedding of an

instance and the embedding of its augmented object should be invariant, and embeddings of different instances should be spread-out. The embedding here is the output of the pre-trained model. In our case, all utterances are divided into segments with equal length, and each segment is taken as a distinct instance in pre-training. For each batch,  $n$  segments are selected with pre-defined sampling strategies which are introduced in Section 2.3. For each segment  $x_i$  in the batch, augmentation is applied and we obtain  $\hat{x}_i$ . Embeddings, denoted by  $f_i$  and  $\hat{f}_i$ , are obtained by feeding the original segment and the augmented segment into the Neural Network (NN) module. To get an augment-invariant embedding, the probability that the augmented instance  $\hat{x}_i$  being classified as instance  $x_i$  is maximized, and is defined as:

$$P(x_i|\hat{x}_i) = \frac{\exp(\langle f_i, \hat{f}_i \rangle / \tau)}{\sum_{k=1}^n \exp(\langle f_k, \hat{f}_i \rangle / \tau)} \quad (1)$$

Embeddings should be instance-spread-out, where the probability that an instance  $x_j$  being classified as another instance  $x_i$ ,  $j \neq i$ , in a batch is minimized, and is defined as:

$$P(x_i|x_j) = \frac{\exp(\langle f_i, f_j \rangle / \tau)}{\sum_{k=1}^n \exp(\langle f_k, f_j \rangle / \tau)}, j \neq i \quad (2)$$

Here, probabilities are calculated as the ratios of exponential inner products of two embeddings scaled by a hyper-parameter  $\tau$ . In all experiments,  $\tau$  is empirically set to be 10.

We assume that the probability of an instance  $x_i$  being recognized as a different instance  $x_j$  is independent for  $j \neq i$ . For each instance  $x_i$ , the joint probability that its augmented instance  $\hat{x}_i$  can be recognized as  $x_i$  and other instances  $x_j$  cannot be recognized as  $x_i$  is denoted as:

$$P_i = P(x_i|\hat{x}_i) \prod_{j \neq i} (1 - P(x_i|x_j)) \quad (3)$$

The NN module is then optimized by minimizing the average negative log-likelihood over a batch. Hence, the loss is :

$$L_{IDL} = -\frac{1}{n} \left( \sum_i \log P(x_i|\hat{x}_i) + \sum_i \sum_{j \neq i} \log(1 - P(x_i|x_j)) \right) \quad (4)$$

## 2.2. Augment-Invariant Embeddings

The first goal of IDL pre-training is to learn augment-invariant embeddings, corresponding to Equation 1. Augmentation techniques are applied to instances in the batch during pre-training. However, for speech signals, effects of augmentation methods on depression status have not been fully explored. To mitigate potential negative effects of some augmentation methods that might change acoustic correlates of depression, such as pitch, formants, etc. [10], augmentation methods used in this work are carefully chosen to be additive noise and volume perturbation at the signal level. For comparison, Vocal Tract Length Perturbation (VTLP) [31], as an augmentation technique that can change formant features, is also applied. All signal-level augmentations are conducted using open-source nlpaug package [32]. In addition, to better analyze the effects of augmentation in the time and frequency domains separately, TM (time masking), FM (frequency masking) and SpecAugment (as a combination of TM and FM along with time-warping) [33] are applied at the feature level, using mel-spectrograms.

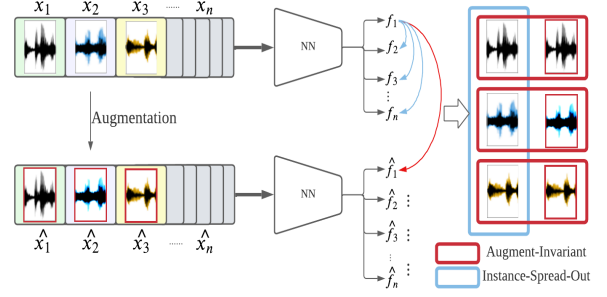


Figure 1: A schematic diagram of IDL pre-training.  $x_i$  indicates an original instance and  $\hat{x}_i$  represents its augmented version,  $i = 1, 2, \dots, n$ , where  $n$  is the batch size.  $f_i$  and  $\hat{f}_i$  denote embeddings of  $x_i$  and  $\hat{x}_i$ . Red and blue blocks are augment-invariant and instance-spread-out computations, respectively.

## 2.3. Instance-Spread-Out Embeddings

### 2.3.1. Distinct speaker-based and Random Sampling

Equation 2 attempts to achieve the goal of training embeddings to be instance-spread-out. Depending on different sampling strategies, the pre-training task has different implicit tendencies. In distinct speaker-based sampling (DS), it is guaranteed that each segment in a batch is from a distinct speaker. Hence, while the model is trying to spread out embeddings from different instances, it is also optimized to classifying different speakers. In other words, speaker information might be preserved in the embedding. The other sampling strategy is random sampling (RS), where segments in a batch are not constrained to be chosen from distinct speakers. If two samples in a batch are from the same speaker, they will still be classified into two instances during pre-training. Therefore, setting the RS strategy might eliminate speaker discriminative information from the embedding.

### 2.3.2. Pseudo Instance-based Sampling

Inspired by Hubert [25], pseudo labels generated by clustering algorithms can reveal implicit correlation between hidden representations and underlying acoustic units in Automatic Speech Recognition (ASR) tasks. Embeddings trained from IDL might also have such a non-trivial correlation with depression status. The correlation can provide distinguishable characteristics across instances which may help with depression classification. Therefore, PIS is proposed to sample instances in a batch according to pseudo labels assigned by clustering algorithms instead of being speaker-dependent. Let  $X$  denote all segments  $X = [x_1, x_2, \dots, x_n]$ . At the first stage, the model is pre-trained using IDL with DS sampling strategy. Then, embeddings,  $F = [f_1, f_2, \dots, f_n]$ , are obtained by feeding  $X$  into the best pre-trained model. Embeddings are clustered using a simple k-means model with  $C$  cluster centroids. Corresponding pseudo labels,  $\hat{Y} = [\hat{y}_1, \hat{y}_1, \dots, \hat{y}_n]$ , are assigned as cluster centroid variables, where  $\hat{y}_i = 1, 2, \dots, C$ . At the second stage, instead of sampling instances in a batch with DS, each instance is sampled from a distinct cluster to guarantee all samples in a batch have distinct pseudo labels. Therefore,  $C$  is pre-determined to be the batch size in the second stage to guarantee all samples in a batch are from different clusters.

Table 1: Performance of IDL in terms of average F1-scores with different augmentation methods using distinct speaker-based sampling (DS) on DAIC-WOZ and CONVERGE test sets. Baseline is the experiments without pre-training. ND and D stands for non-depression and depression, respectively. In the Augmentation column, TM stands for Time-Masking, FM is Frequency-Masking, and SpecAug stands for SpecAugment. VTLP denotes Vocal Tract Length Perturbation. Noise and Volume stands for noise perturbation and volume perturbation, respectively. The best F1-scores are boldfaced. \* indicates that the change is not statistically significant.

Exp	Augmentation	DAIC-WOZ			CONVERGE		
		F1 (avg)	F1(ND)	F1(D)	F1 (avg)	F1(ND)	F1(D)
Baseline	NA	0.4258	0.5357	0.3158	0.7228	0.7101	0.7355
	TM	<b>0.6458</b>	0.8116	0.4800	<b>0.7412*</b>	0.7503	0.7321
IDL - Feature Level Aug	FM	0.6103	0.7761	0.4444	0.7256	0.7343	0.7168
	SpecAug	0.6189	0.8378	0.4000	0.7316	0.7432	0.7200
IDL - Signal Level Aug	VTLP	0.5428	0.6984	0.3871	0.7197	0.7378	0.7015
	Noise	0.6083	0.8000	0.4167	0.7293	0.7537	0.7048
	Volume	<b>0.6458</b>	0.8116	0.4800	0.7257	0.7453	0.7063

### 3. Experimental Setup

For both pre-training and downstream datasets, we use 40 dimensional mel-spectrograms as input features, extracted with the Librosa library [34] every 32ms with a 64ms Hanning window. Experiments are conducted using PyTorch [35].

#### 3.1. Datasets

##### 3.1.1. DAIC-WOZ

The Distress Analysis Interview Corpus - Wizard of Oz (DAIC-WOZ) [36] is an English dataset containing audio, text and video of interviews from 189 male and female participants, collected by a virtual interviewer. Audio files range between 7-33 min (16 min on average), with a non-depression vs depression ratio of 3:1. In this work, only recordings belonging to participants are extracted according to the groundtruth time stamps. Train, validation and test sets are partitioned in the same way as that provided in the dataset [36].

Librispeech [37] is used as the pre-training dataset for downstream tasks on DAIC-WOZ. It is the largest publicly available speech corpus in English with mostly reading style speech. Two subsets, training and validation, are partitioned with the ratio of 9:1 without speaker overlap, and the validation set is used to choose the best pre-trained model.

##### 3.1.2. CONVERGE

The Mandarin dataset we use is a part of the China, Oxford and Virginia Commonwealth University Experimental Research Genetic Epidemiology (CONVERGE) study [38]. The dataset contains 391 hours of recordings from 7959 females with almost equal numbers of depression and non-depression classes. The sampling rate was 16kHz, and train, validation and test sets are split into 60%, 20% and 20%.

CN-Celeb [39] is chosen as the pre-training dataset for classification experiments conducted on CONVERGE. The dataset contains more than 130,000 utterances in Mandarin from 1000 Chinese celebrities across multiple genres. Training and validation subsets are partitioned in the same way as Librispeech.

#### 3.2. DepAudioNet

The back-end model we apply is DepAudioNet proposed in [16, 20] for depression detection from speech signals. 40-dimensional Mel-spectrogram features first pass through a one-dimensional convolution layer with a kernel size of 3 to capture short-term characteristics. Batch normalization is enabled after

convolution. A max-pooling layer with a kernel size of 3 follows the convolution layer to capture mid-term features, with a dropout factor of 0.05 and an activation function as ReLU. Two LSTM layers and a fully connected layer activated by a Sigmoid with a hidden size of 128 are concatenated to generate binary predictions.

For experiments on DAIC-WOZ, the model is pre-trained on Librispeech corpus. Utterances from each speaker are combined together and segmented into multiple fixed length segments with 120 frames each. The model is trained for 100 epochs with a batch size of 20. The learning rate is set to 1e-3, and the decay factor is set to 0.9 every two epochs. The model with the smallest validation loss is chosen to initialize the downstream model. For the depression classification task, to mitigate length variation and class imbalance, random cropping and random sub-sampling are applied [16]. Each utterance is randomly cropped into a fragment with the same length as the shortest utterance, to mitigate any influence from longer objects. Then, each fragment is segmented with a fixed window length of 120 frames. A new subset is generated by sampling an equal number of depression and non-depression segments randomly without replacement. To fully utilize as many samples as possible, five individual models are trained using five randomly selected subsets for 100 epochs each, with a learning rate of 1e-3, and a decay factor of 0.9 every two epochs. Final predictions are obtained by averaging probabilities of the five models.

For CONVERGE experiments pre-trained on CN-Celeb, random cropping and sub-sampling are disabled because the dataset is balanced. One model is trained using all segments. The learning rate is empirically set to 1e-2. Other experimental configurations are the same as DAIC-WOZ experiments.

## 4. Results and Discussion

#### 4.1. Comparison of Different Augmentation Methods

To investigate the effects of augment-invariant characteristics, different augmentation methods used in pre-training are compared while fixing the sampling strategy to be DS. F1-scores for the baseline system without pre-training, and IDL pre-training with various augmentation methods on DAIC-WOZ and CONVERGE test sets are shown in Table 1. All best F1-scores are statistically significant compared with the baseline unless specified. In the feature-level augmentation experiments, the effect of time and frequency perturbation on depression status can be analyzed by perturbing the spectrogram in time or frequency at a time. Table 1 shows that TM gives the best performance with

a relative average F1-Score improvement of 51.67% on DAIC-WOZ and 2.55% on CONVERGE compared with the baseline system. FM is the worst among the three with a relative improvement of 43.33% and 0.39% on DAIC-WOZ and CONVERGE, respectively. Applying SpecAugment achieves an intermediate performance with a relative improvement of 45.35% and 1.22%, respectively. The results suggest that, augmenting the signal in the spectral domain, such as frequency masking, might result in loss of depression-specific information. The worst performance is observed for IDL with VTLP as the augmentation method, which also proves the hypothesis that modifications of spectral domain parameters can negatively affect depression classification. Because the spectrum is less affected by noise and volume perturbations compared with VTLP, moderate improvement can be observed on both datasets using these perturbations. Combination of augmentation methods are evaluated but performances are worse than using a single method.

#### 4.2. Sampling Strategy Comparison

The instance-spread-out characteristics of the proposed pre-training method are explored by setting different sampling strategies with TM augmentation, since TM is the best augmentation technique found in the experiments reported in Section 4.1. Average F1-scores on DAIC-WOZ and CONVERGE test sets using different sampling strategies are shown in Table 2.

In Table 2, we observe that setting the batch sampling strategy to be DS yields relative improvements of 13.64% and 4.79% on DAIC-WOZ and CONVERGE, respectively, compared with using RS. Unlike RS, DS might preserve speaker information because embeddings of two speakers’ instances are optimized to be spread-out, this observation implies speaker-discriminative information might be crucial in determining depression status. An ablation study is conducted in Section 4.3 to prove that speaker information is preserved using DS.

Table 2: Average F1-scores of Baseline and unsupervised pre-training methods on DAIC-WOZ and CONVERGE. SS stands for sampling strategy. Best Results are boldfaced.

F1(avg)	SS	DAIC-WOZ	CONVERGE
Baseline	-	0.4258	0.7228
CPC[40]	-	0.5104	0.7371
Speech SimCLR[28]	-	0.5747	0.7288
IDL	RS	0.5683	0.7073
	DS	0.6458	0.7412
	PIS	<b>0.6834</b>	<b>0.7435</b>

In the IDL-PIS experiments, TM is also chosen for augmentation. PIS can further improve depression classification performance compared with DS, with relative improvements of 5.82% on DAIC-WOZ and 0.31% on CONVERGE. Improvements demonstrate that pseudo-labels generated by the clustering model provide a high correlation with depression status.

As a comparison with other UPT methods, experiments with CPC [40] and Speech SimCLR [28] without reconstruction loss are conducted and reported in Table 2. Results show that CPC and Speech SimCLR can perform better than the baseline system but not as well as the proposed IDL method.

#### 4.3. Ablation Study on Speaker Classification

We have shown that sampling instances in a batch using the DS sampling strategy during pre-training can help with the down-

Table 3: Average F1-scores and Speaker Classification Accuracy (Spk Cls Acc) of Baseline and unsupervised pre-training on DAIC-WOZ. w/o ft stands for no fine-tuning

Exp	F1(avg)	Spk Cls Acc(%)
Baseline	0.4258	21.98
IDL (DS w/o ft)	0.3472	68.26
IDL (DS )	0.6458	72.16
IDL (RS)	0.5683	59.45
IDL (PIS)	0.6834	71.60

stream depression classification task. To prove that speaker information is preserved in embeddings optimized by pre-training, speaker classification tasks with embeddings generated from downstream models as input are conducted on the DAIC-WOZ test set using a simple SVM classifier [41]. 30% of the segments constitute the speaker classification test set.

Table 3 shows that, for IDL pre-training using TM with DS on Librispeech, without fine-tuning on the DAIC-WOZ, a 68.26% speaker classification accuracy is achieved. The corresponding F1-score of 0.3472 is reasonable since the model hasn’t been tuned for the depression task. The accuracy increases to 72.16% after fine-tuning. This improvement can be explained by in-domain dataset adaptation through the downstream task. As a comparison, IDL (RS) only achieves 59.45% speaker classification accuracy. The relative speaker classification accuracy degradation of 17.61% from DS to RS proves that speaker information are better preserved using DS. Additionally, using PIS achieves a comparable speaker classification accuracy and an improved F1-score on depression classification compared with DS. This observation reveals that depression-specific characteristics can be preserved in embeddings trained using PIS, along with some speaker information.

## 5. Conclusion and Future Work

In this paper, a modified UPT approach, IDL, is proposed to learn augment-invariant and instance-spread-out embeddings for depression detection tasks on DAIC-WOZ and CONVERGE datasets. Different augmentation techniques are compared in terms of augment-invariant characteristics. Results show that TM yields the best performance among all augmentation methods. For learning instance-spread-out embeddings, different sampling strategies, DS and RS are investigated and compared. Results show that preserving speaker information in the embedding using DS might help with depression classification. We also propose a new sampling strategy, PIS, to generate pseudo labels based on clustering, to reveal a deeper correlation between embeddings and depression status. Compared with the baseline without pre-training, the proposed approach, PIS, achieves significant improvements in the detection of MDD on DAIC-WOZ and CONVERGE datasets, with  $p$ -value 0.0015 and 0.05, respectively. In future work, we will investigate context-based, in addition to speaker-based, sampling strategies and apply IDL and PIS to other low-resource tasks.

## 6. Acknowledgements

This work was funded in part 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] M. A. Lopez Molina, K. Jansen, C. Drews, R. Pinheiro, R. Silva, and L. Souza, "Major depressive disorder symptoms in male and female young adults," *Psychology, Health & Medicine*, vol. 19, no. 2, pp. 136–145, 2014.
- [2] W. H. Organization *et al.*, "Depression and other common mental disorders: global health estimates," World Health Organization, Tech. Rep., 2017.
- [3] C. Otte, S. M. Gold *et al.*, "Major depressive disorder," *Nature Reviews Disease Primers*, vol. 2, no. 1, pp. 1–20, 2016.
- [4] K. Kroenke, T. 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.
- [5] D. M. Khan, N. Yahya, N. Kamel, and I. Faye, "Automated diagnosis of major depressive disorder using brain effective connectivity and 3d convolutional neural network," *IEEE Access*, vol. 9, pp. 8835–8846, 2021.
- [6] Z. Huang, J. Epps, and D. Joachim, "Investigation of speech landmark patterns for depression detection," *IEEE Transactions on Affective Computing*, 2019.
- [7] A. Pampouchidou, K. Marias, M. Tsiknakis, P. Simos, F. Yang, and F. Meriaudeau, "Designing a framework for assisting depression severity assessment from facial image analysis," in *ICSIPA*, 2015, pp. 578–583.
- [8] L. Wen, X. Li, G. Guo, and Y. Zhu, "Automated depression diagnosis based on facial dynamic analysis and sparse coding," *IEEE Transactions on Information Forensics and Security*, vol. 10, no. 7, pp. 1432–1441, 2015.
- [9] U. R. Acharya, V. K. Sudarshan, H. Adeli, J. Santhosh, J. E. Koh, and A. Adeli, "Computer-aided diagnosis of depression using eeg signals," *European Neurology*, vol. 73, no. 5-6, pp. 329–336, 2015.
- [10] N. Cummins, S. Scherer *et al.*, "A review of depression and suicide risk assessment using speech analysis," *Speech Communication*, vol. 71, pp. 10–49, 2015.
- [11] P. Lopez-Otero, L. Docio-Fernandez, and C. Garcia-Mateo, "A study of acoustic features for the classification of depressed speech," in *MIPRO*. IEEE, 2014, pp. 1331–1335.
- [12] Z. Liu, B. Hu *et al.*, "Detection of depression in speech," in *2015 ACII*. IEEE, 2015, pp. 743–747.
- [13] L.-S. A. Low, N. C. Maddage, M. Lech, L. Sheeber, and N. Allen, "Influence of acoustic low-level descriptors in the detection of clinical depression in adolescents," in *ICASSP*, 2010, pp. 5154–5157.
- [14] E. Rejaibi, A. Komaty, F. Meriaudeau, S. Agrebi, and A. Othmani, "Mfcc-based recurrent neural network for automatic clinical depression recognition and assessment from speech," *Biomedical Signal Processing and Control*, vol. 71, p. 103107, 2022.
- [15] U. Sharma, "Detection of depression from speech signal through linear svm using mfcc feature," in *ICICC*, 2020.
- [16] X. Ma, H. 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.
- [17] Z. Huang, J. Epps, and D. Joachim, "Exploiting vocal tract coordination using dilated cnns for depression detection in naturalistic environments," in *ICASSP*, 2020, pp. 6549–6553.
- [18] A. Afshan, J. Guo *et al.*, "Effectiveness of voice quality features in detecting depression," *Interspeech*, 2018.
- [19] T. F. Quatieri and N. Malyska, "Vocal-source biomarkers for depression: A link to psychomotor activity," in *Thirteenth Annual Conference of ISCA*, 2012.
- [20] A. Bailey and M. D. Plumbley, "Raw audio for depression detection can be more robust against gender imbalance than mel-spectrogram features," *arXiv preprint arXiv:2010.15120*, 2020.
- [21] V. Ravi, J. Wang, J. Flint, and A. Alwan, "Fraug: A frame rate based data augmentation method for depression detection from speech signals," *arXiv preprint arXiv:2202.05912*, 2022.
- [22] A. Salekin, J. W. Eberle, J. J. Glenn, B. A. Teachman, and J. A. Stankovic, "A weakly supervised learning framework for detecting social anxiety and depression," *IMWUT*, vol. 2, no. 2, pp. 1–26, 2018.
- [23] P. Zhang, M. Wu, H. Dinkel, and K. Yu, "Depa: Self-supervised audio embedding for depression detection," in *Proceedings of the 29th ACM-MM*, 2021, pp. 135–143.
- [24] A. Vázquez-Romero and A. Gallardo-Antolín, "Automatic detection of depression in speech using ensemble convolutional neural networks," *Entropy*, vol. 22, no. 6, p. 688, 2020.
- [25] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM TASLP*, vol. 29, pp. 3451–3460, 2021.
- [26] R. Fan, A. Afshan, and A. Alwan, "Bi-apc: Bidirectional autoregressive predictive coding for unsupervised pre-training and its application to children's asr," in *ICASSP*, 2021, pp. 7023–7027.
- [27] M. Li, B. Yang, J. Levy, A. Stolcke, V. Rozgic, S. Matsoukas, C. Papayannis, D. Bone, and C. Wang, "Contrastive unsupervised learning for speech emotion recognition," in *ICASSP*, 2021, pp. 6329–6333.
- [28] D. Jiang, W. Li, M. Cao, W. Zou, and X. Li, "Speech simplr: Combining contrastive and reconstruction objective for self-supervised speech representation learning," *arXiv preprint arXiv:2010.13991*, 2020.
- [29] S. Sardari, B. Nakisa, M. N. Rastgoo, and P. Eklund, "Audio based depression detection using convolutional autoencoder," *Expert Systems with Applications*, vol. 189, p. 116076, 2022.
- [30] M. Ye, X. Zhang, P. C. Yuen, and S.-F. Chang, "Unsupervised embedding learning via invariant and spreading instance feature," in *Proceedings of the IEEE/CVF CVPR*, 2019, pp. 6210–6219.
- [31] N. Jaitly and G. E. Hinton, "Vocal tract length perturbation (vtlp) improves speech recognition," in *Proc. ICML Workshop on Deep Learning for Audio, Speech and Language*, vol. 117, 2013, p. 21.
- [32] E. Ma, "Nlp augmentation," <https://github.com/makcedward/nlpaug>, 2019.
- [33] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," *arXiv preprint arXiv:1904.08779*, 2019.
- [34] B. McFee, C. Raffel, D. Liang, D. P. Ellis, M. McVicar, E. Battenberg, and O. Nieto, "Librosa: Audio and music signal analysis in python," in *Proceedings of the 14th Python in Science Conference*, vol. 8. Citeseer, 2015, pp. 18–25.
- [35] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library," *Advances in NIPS*, vol. 32, pp. 8026–8037, 2019.
- [36] M. Valstar, J. 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.
- [37] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *ICASSP*, 2015, pp. 5206–5210.
- [38] Y. 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.
- [39] Y. Fan, J. Kang *et al.*, "Cn-celeb: a challenging chinese speaker recognition dataset," in *ICASSP*, 2020, pp. 7604–7608.
- [40] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [41] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.