



Published in final edited form as:

Interspeech. 2022 September ; 2022: 3338–3342. doi:10.21437/interspeech.2022-10798.

A Step Towards Preserving Speakers' Identity While Detecting Depression Via Speaker Disentanglement

Vijay Ravi[†], Jinhan Wang[†], Jonathan Flint[§], Abeer Alwan[†]

[†]Dept. of Electrical and Computer Engineering, University of California, Los Angeles, USA

[§]Dept. of Psychiatry and Biobehavioral Sciences, University of California, Los Angeles, USA

Abstract

Preserving a patient's identity is a challenge for automatic, speech-based diagnosis of mental health disorders. In this paper, we address this issue by proposing adversarial disentanglement of depression characteristics and speaker identity. The model used for depression classification is trained in a speaker-identity-invariant manner by minimizing depression prediction loss and maximizing speaker prediction loss during training. The effectiveness of the proposed method is demonstrated on two datasets - DAIC-WOZ (English) and CONVERGE (Mandarin), with three feature sets (Mel-spectrograms, raw-audio signals, and the last-hidden-state of Wav2vec2.0), using a modified DepAudioNet model. With adversarial training, depression classification improves for every feature when compared to the baseline. Wav2vec2.0 features with adversarial learning resulted in the best performance (F1-score of 69.2% for DAIC-WOZ and 91.5% for CONVERGE). Analysis of the class-separability measure (J-ratio) of the hidden states of the DepAudioNet model shows that when adversarial learning is applied, the backend model loses some speaker-discriminability while it improves depression-discriminability. These results indicate that there are some components of speaker identity that may not be useful for depression detection and minimizing their effects provides a more accurate diagnosis of the underlying disorder and can safeguard a speaker's identity.

Index Terms:

paralinguistics; depression detection; adversarial learning; speaker disentanglement; privacy in healthcare

1. Introduction

Mental health disorders (e.g. Major depressive disorder or MDD) are identifiable by verbal cues such as monotonic-speech, choice of vocabulary, abnormal disfluencies, etc [1, 2]. Previous studies have identified discernible differences between speech patterns of depressed and healthy subjects [3, 4]. In addition, collecting speech data has become increasingly easier with the advent of digital voice assistants. As a result, speech-based automatic detection of MDD has received special attention in recent years [5, 6, 7].

The effectiveness of speech features - spectral [8, 9], prosodic [10], articulatory [11] and voice-quality [12] - for classification of MDD, have been analyzed in the past. Lately, deep-learning methods for MDD detection have become popular as they outperform traditional pattern recognition techniques [13, 14, 15, 16, 17].

Among others, acoustic features such as x-vectors [18], i-vectors [19] and other speaker embeddings [20] have been shown to be effective in the diagnosis of a speaker's mental state. These features, however, also carry information about a speaker's identity [21] which can be counter-productive to privacy preservation- a key factor in the adoption of digital mental-health screening systems [22]. Subsequently, an important question that has remained unanswered in the speech-research community is whether depression detection can be done in a speaker-identity-invariant manner. Additionally, it is still not known if there are components of speech characterizing a speaker that may not be relevant to their mental health status. More recently, two studies introduced algorithms to preserve privacy during depression detection; [23] proposed sine-wave speech representation and [24] used federated learning. However, the performance of depression detection in both studies degrades while preserving patient's privacy. In this paper, the paradigm of adversarial learning to disentangle speaker and depression characteristics is investigated.

In the past, adversarial speaker normalization has been evaluated in the domain of emotion recognition [25, 26, 27]. In [25], the authors perform speaker-invariant domain adaptation on multi-modal features (speech, text, and video) for emotion recognition. In [26], gradient reversal technique with an entropy loss is proposed to disentangle emotion and speaker information. In [27], the authors fine-tune a pre-trained Hubert-base model ([28], 300M parameters) with gradient-based adversarial learning. Fine-tuning such models can require large amounts of in-domain data and be computationally intensive. Moreover, these papers utilize IEMOCAP and MSP-Improv datasets which are mono-lingual and consists of acted audio data [29, 30].

In contrast, in this paper, we propose adversarial disentanglement of speaker-identity and depression information, using speech-features only, on datasets with continuous and non-acted speech. In addition to Mel-spectrograms and raw-audio signals, we propose the use of the last-hidden-state of a pre-trained Wav2Vec2.0 model [31] as the input feature. Wav2vec2.0 models are trained on large amounts of unlabelled data making them robust in various speech processing tasks. The Wav2vec2.0 model is used as a feature extractor, thereby reducing fine-tuning computations.

Further, unlike prior studies in emotion recognition, we show that the benefits of the proposed method extend to another language (Mandarin). Lastly, we analyze the class separability power of the hidden states of the backend model. We find that while speaker-separability is reduced with adversarial learning, depression-discrimination capability improved.

To the best of our knowledge, this is the first work to address privacy in depression diagnosis from a speaker-disentanglement perspective. We hypothesize that not all speaker information is useful for depression detection and successfully demonstrate that normalizing irrelevant

speaker-related information can result in a more accurate diagnoses, not to mention help protect patients' identity.

2. Adversarial Learning

We propose a loss-based adversarial learning mechanism for speaker-disentangled depression detection. Inspired from the domain-adversarial training proposed in [32], our approach involves a loss minimization-maximization technique.

Let the number of unique speakers in the training data be N . The loss used for the prediction of MDD binary labels is:

$$L_{MDD} = -\frac{1}{N} \sum_{i=1}^N [Y_i \log(p_i) + (1 - Y_i) * \log(1 - p_i)] \quad (1)$$

$Y_i \in (0, 1)$ is the class label for the i^{th} speaker and p_i is the probability that speaker i is depressed. The loss for speaker ID prediction is defined as -

$$L_{SPK} = -\frac{1}{N} \sum_{n=1}^N \left[\log \frac{\exp(x_{n,n})}{\sum_{c=1}^N \exp(x_{n,c})} \right], \quad (2)$$

where $x_{n,n}$ is the output score that the n^{th} speaker is predicted correctly and $x_{n,c}$ is the score of the n^{th} speaker being predicted as another speaker c .

To train the model in a speaker-identity-invariant manner, during optimization, we minimize the depression loss and maximize the speaker prediction loss. This can be written as:

$$L_{total} = L_{MDD} - \lambda(L_{SPK}), \quad (3)$$

where λ is an empirically determined hyperparameter that controls how much of the speaker loss contributes to the total loss. Initial lambda values were selected to be similar to those reported in the literature (1e-3) [25]. We experimented with higher and lower values and chose the best performing lambda values. By this process, we force the model to focus more on depression-discriminatory information and ignore some speaker-discriminatory information, thereby making the model invariant to changes in some speaker-specific characteristics.

3. Experimental Details

Experiments were conducted using the DAIC-WOZ (English, [33]) and CONVERGE (Mandarin, [34]) datasets. The backend model was a modified version of the DepAudioNet [13] which was trained using Mel-spectrograms, raw-audio signals and the last-hidden-state of Wav2vec2.0 as input features. Details of the datasets, model parameters, input features, and evaluation metrics used are described in this section.

3.1. Datasets

3.1.1. DAIC-WOZ—The Distress analysis interview corpus wizard of Oz (DAIC-WOZ) [33] 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 [35]. 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 training because of time-labelling errors. This dataset consists of 58 hours of audio data, sampled at 16kHz. Data partitioning is the same as that provided with the database description (107 speakers for training and 35 speaker for evaluation).

3.1.2. CONVERGE—The CONVERGE (China, Oxford and Virginia Commonwealth University Experimental Research on Genetic Epidemiology) dataset [34] comprises of audio interviews of 7959 female participants being interviewed by a trained interviewer. This study focused on subjects with increased genetic risk for MDD, and to obtain a more genetically homogeneous sample, only women participants were recruited. The diagnoses of depressive disorders were made with the Composite International Diagnostic Interview (Chinese version) [36]. Audio recordings are sampled at a rate of 16kHz. In this study, a subset of this dataset is used with 1185 randomly chosen speakers in the training partition and the remaining 507 speakers in the evaluation partition. This database is characterized by a large degree of phonetic and content variability.

3.2. Model - DepAudioNet

DepAudioNet was chosen as a baseline because its code is publicly available in addition to its high accuracy of depression classification. Model implementation was based on [37] and network parameters such as the number of hidden layers, learning rate, and dropout probability etc. were chosen empirically. Kernel size and stride of the *Conv1D* layer are denoted by K and S , respectively. Hidden state dimension of the recurrent LSTM layers is denoted by H .

For the DAIC-WOZ dataset, the model with Mel-spectrograms as input consisted of one *Conv1D* layer ($K = 3$, $S = 1$) and four unidirectional LSTM layers ($H = 128$). In case of raw-audio signals, two *Conv1D* layers ($K_1 = 1024$, $S_1 = 512$, $K_2 = 3$, $S_2 = 1$) and two LSTM layers ($H = 128$) were used. For Wav2vec2.0 features, the model consisted of an input LSTM layer ($H = 256$) followed by five hidden LSTM layers with the same H as the input layer. For the CONVERGE data and Mel-spectrograms, the model consisted of two *Conv1D* layers ($K_1 = 3$, $S_1 = 1$, $K_2 = 3$, $S_2 = 1$) and four unidirectional LSTM layers ($H = 128$). In case of CONVERGE and raw-audio signals, two *Conv1D* layers ($K_1 = 1024$, $S_1 = 512$, $K_2 = 3$, $S_2 = 1$) and four LSTM layers ($H = 512$) were used. For Wav2vec2.0 features with CONVERGE data, the model consisted of an input LSTM layer ($H = 256$) followed by five hidden LSTM layers with the same H as the input layer.

Conv1D layers were followed by ReLU non-linearity, a dropout layer and a max-pooling layer with a kernel of size 3. For every model configuration, the final layers were fully connected layers to generate the predictions for MDD and speaker labels. Based on the

number of speakers in the training set, output dimensions for speaker labels were 107 for experiments with DAIC-WOZ and 1185 for CONVERGE. For MDD prediction, a sigmoid activation was used with binary cross entropy loss, while for speaker ID prediction, cross entropy loss was used without any output activation.

3.3. Input Features

Three feature sets were analyzed – 1) 40-dimensional Mel-spectrograms extracted using a Hanning window of length $w = 1024$ samples (64ms), and hop size $h = 512$ samples (32ms), 2) raw-audio signal, and 3) the last-hidden-state of Wav2vec2.0 model [31]. For the DAIC-WOZ data, Wav2vec2.0 base model was used where the hidden state dimension was 768. For CONVERGE, Wav2vec2.0 large model (XLSR-Mandarin) was used where the hidden state dimension was 1024 [38]. Mean-variance normalization was applied to Mel-spec and raw-audio features. Wav2vec2.0-base and XLSR-Mandarin models from the Hugging Face library [39, 40] were used.

In the case of the DAIC-WOZ dataset, training data were pre-processed by random cropping and sampling [13]. Each utterance was randomly cropped to fragments (equal to the length of the shortest utterance) and every fragment was segmented into multiple segments. Segment lengths were 3.84s each which translates into 120 frames for Mel-spectrogram, 61440 samples for raw-audio and 200 frames for Wav2vec2.0 features. A training subset was generated by randomly sampling, without replacement, an equal number of depression and non-depression segments. In each experiment, five separate models were trained using a randomly generated training subset and the final prediction was averaged across those five models. On the contrary, for the CONVERGE dataset, segments were generated without random cropping and sampling with the segment length the same as before, and each experiment was performed by training one model using all of the training data.

3.4. Evaluation Metrics

Depression classification was evaluated using the F1-score to account for both false positives and false negatives [41].

The J-ratio is a metric used to measure class separability of a given set of features [42] and is computed as follows:

$$S_W = \frac{1}{N} \sum_{i=1}^N R_i, \quad (4)$$

$$S_B = \frac{1}{n} \sum_{i=1}^N (M_i - M_o)(M_i - M_o)^T, \quad (5)$$

$$J = \text{trace}[(S_W + S_B)^{-1} S_B], \quad (6)$$

where N is the total number of speakers, R_i is the covariance matrix for the i^{th} speaker, M_i is the mean vector for the i^{th} speaker, and M_o is the mean of all M_i s. S_W and S_B are the

within- and between-class scatter matrices, respectively. Hidden states of all recurrent layers (LSTMs) of the DepAudioNet model were extracted and used as inputs in the calculation of J-ratios. This metric was used to examine speaker separability [43].

4. Results and Discussion

Results are presented and discussed in two steps. First, performance of the proposed method relative to the baseline system is presented for each of the three input features and the two datasets. Then, through speaker-separability analysis, it is shown that the proposed method indeed normalizes some speaker information resulting in better MDD classification. Relative improvements, unless specified, are statistically significant [44].

4.1. Adversarial Learning

Table 1 shows the performance of the proposed approach using the DAIC-WOZ dataset in terms of the F1-score for the Non-Depressed Class (ND), the Depressed class (D) and the average F1-score (F1-Avg) across both classes. In the case of Mel-spectrograms, the baseline DepAudioNet model [13] has an F1-Avg of 0.619 which is in agreement with previous work [37]. When adversarial speaker-invariant training is applied, the overall relative performance improvement is 4.36% where F1-ND increased by 3.6% and F1-D by 5.06%. The baseline performance with raw audio signals as input is better than the Mel-spec baseline. In this case, with adversarial learning, while F1-ND is reduced by 6.8%, F1-D increased by 16% resulting in an overall relative improvement of 2.16%. With Wav2vec2.0 features, in spite of a very high baseline performance (F1-Avg = 0.686), the proposed method provides overall improvements of 0.87% resulting in the best F1-Avg score of 0.692.

Similar improvements are observed using the CONVERGE dataset. As seen in Table 2, for Mel-spectrograms, relative performance increases by 1.25%, from 0.879 for the baseline to 0.890 for the proposed method. Improvements for raw-audio are higher; baseline system has an F1-Avg of 0.829 whereas the proposed method has an F1-Avg of 0.857 (relative improvement of 3.34%). Lastly, for the features extracted from the XLSR-Mandarin model, a relative improvement of 0.33% (not statistically significant) is observed - F1-score for baseline is 0.912 and 0.915 for the proposed adversarial disentanglement, respectively. For both datasets, features extracted from Wav2vec2.0 models, which are trained in an unsupervised manner and fine-tuned for speech recognition, performed the best.

Another important observation to note is that adversarial disentanglement of speaker and depression was sensitive to λ values and in most experiments, a very small value resulted in significant improvements.

Relative improvements in system performance are summarized in Table 3. The proposed method improves depression prediction score for every feature across two different datasets. More importantly, prediction performance improves for both ND and D classes (except for DAIC-WOZ raw-audio), demonstrating that the proposed method can improve precision and recall of automatic MDD diagnosis in most cases. Improvements, however, are marginal for content-rich features such as Wav2vec2.0 base and XLSR-Mandarin. It maybe that since

these model are fine-tuned for speech recognition tasks, much of the speaker-information is already lost during feature extraction resulting in smaller improvements.

4.2. Speaker Separability Analysis

Speaker separability, measured in terms of J-ratio, along with F1-Avg scores for MDD prediction are presented in Tables 4 and 5 for the DAIC-WOZ and CONVERGE datasets, respectively. A higher J-ratio indicates better speaker discrimination. For the DAIC-WOZ dataset, the proposed method reduces J-ratio for every feature while improving the F1-Score. Interestingly, configurations with lower speaker separability performed better in depression classification - for example baseline systems vs. adversarial disentanglement and Wav2vec2.0 vs. Mel-spectrogram.

Contrary to DAIC-WOZ, the J-ratio for XLSR-Mandarin using the CONVERGE data was higher than those for Mel-spec. It is possible that Wav2vec2.0 captures room acoustics resulting in a higher speaker separability. Nevertheless, trends are similar, in that the J-ratio for speaker separation reduces for every feature set when the proposed method is applied. Comparing the improvements from Table 3 with J-ratios presented in Tables 4 and 5, it can be seen that performance gains are maximum for features with the highest baseline speaker J-ratio; Mel-spec in DAIC-WOZ (+4.36% improvement and baseline J-ratio of 4.81) and Raw-audio in CONVERGE (+3.34% improvement and baseline J-ratio of 302.3). These results confirm our hypothesis that speaker-identity contains some components irrelevant to the mental state of that person and training the model to be invariant to such characteristics can lead to better depression detection.

5. Conclusion and Future Work

In the past, features such as x-vectors, i-vectors, etc. have been shown to be useful for depression detection. These features, despite their effectiveness, contain information about the speaker's identity. Excessive dependence on speaker-identity features can harm the privacy factor of an MDD diagnosis system, an important attribute towards the adoption of speech-based assessment methods. Naturally, a question that arises is whether depression detection can be done in a speaker-invariant manner.

In this paper, we attempt to address this important question by proposing adversarial disentanglement of speaker-identity and depression status. The proposed method illustrates that speaker-identity invariant models can provide better MDD classification performance across multiple features and on multi-lingual datasets. The results presented in this paper support the hypothesis that when correlates of speaker identity, irrelevant to a subject's mental state, are partially normalized, depression diagnosis is more accurate.

Future work will investigate the specifics of speaker-related information that are or are not useful for depression detection, and will compare the effectiveness of content-rich vs. speaker-information-rich features. Other metrics for speaker recognition, such as multi-class accuracy, will also be analyzed.

Acknowledgments

This work was funded in part by NIH award number R01MH122569.

References

- [1]. Nilsson A, "Speech characteristics as indicators of depressive illness," *Acta Psychiatrica Scandinavica*, vol. 77, no. 3, pp. 253–263, 1988. [PubMed: 3394527]
- [2]. Andreasen NJ et al. , "Linguistic analysis of speech in affective disorders," *Archives of General Psychiatry*, vol. 33, no. 11, pp. 1361–1367, 1976. [PubMed: 985047]
- [3]. Cummins N et al. , "A review of depression and suicide risk assessment using speech analysis," *Speech Communication*, vol. 71, pp. 10–49, 2015.
- [4]. France DJ et al. , "Acoustical properties of speech as indicators of depression and suicidal risk," *IEEE transactions on Biomedical Engineering*, vol. 47, no. 7, pp. 829–837, 2000. [PubMed: 10916253]
- [5]. Alghowinem S et al., "Detecting depression: a comparison between spontaneous and read speech," in *ICASSP. IEEE*, 2013, pp. 7547–7551.
- [6]. Ringeval F et al. , "Avec 2019 workshop and challenge: state-of-mind, detecting depression with ai, and cross-cultural affect recognition," in *Proceedings of the 9th AVEC*, 2019.
- [7]. Low DM et al. , "Automated assessment of psychiatric disorders using speech: A systematic review," *Laryngoscope Investigative Otolaryngology*, vol. 5, no. 1, pp. 96–116, 2020. [PubMed: 32128436]
- [8]. Sanchez MH et al. , "Using prosodic and spectral features in detecting depression in elderly males," in *Interspeech*, 2011, pp. 3001–3004.
- [9]. Dubagunta SP et al., "Learning voice source related information for depression detection," in *ICASSP. IEEE*, 2019, pp. 6525–6529.
- [10]. Yang Y et al. , "Detecting depression severity from vocal prosody," *IEEE transactions on affective computing*, vol. 4, no. 2, pp. 142–150, 2012.
- [11]. Seneviratne N et al. , "Multimodal depression classification using articulatory coordination features and hierarchical attention based text embeddings," *arXiv preprint arXiv:2202.06238*, 2022.
- [12]. Afshan A et al. , "Effectiveness of voice quality features in detecting depression," *Interspeech*, 2018.
- [13]. Ma X 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.
- [14]. Rejaibi E et al. , "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]. Shen Y et al. , "Automatic depression detection: An emotional audio-textual corpus and a gru/bilstm-based model," *arXiv preprint arXiv:2202.08210*, 2022.
- [16]. Chlasta K et al. , "Automated speech-based screening of depression using deep convolutional neural networks," *Procedia Computer Science*, vol. 164, pp. 618–628, 2019.
- [17]. Harati A et al., "Speech-based depression prediction using encoder-weight-only transfer learning and a large corpus," in *ICASSP. IEEE*, 2021, pp. 7273–7277.
- [18]. Ravi V et al. , "Frag: A frame rate based data augmentation method for depression detection from speech signals," *arXiv preprint arXiv:2202.05912*, 2022.
- [19]. Di Y et al. , "Using i-vectors from voice features to identify major depressive disorder," *Journal of Affective Disorders*, vol. 288, pp. 161–166, 2021. [PubMed: 33895418]
- [20]. Dumpala SH et al. , "Significance of speaker embeddings and temporal context for depression detection," *arXiv preprint arXiv:2107.13969*, 2021.
- [21]. Snyder D et al., "X-vectors: Robust dnn embeddings for speaker recognition," in *ICASSP. IEEE*, 2018, pp. 5329–5333.

- [22]. Lustgarten SD et al. , “Digital privacy in mental healthcare: current issues and recommendations for technology use,” *Current opinion in psychology*, vol. 36, pp. 25–31, 2020. [PubMed: 32361651]
- [23]. Dumpala SH et al. , “Sine-Wave Speech and Privacy-Preserving Depression Detection,” in Proc. SMM21, Workshop on Speech, Music and Mind 2021, 2021, pp. 11–15.
- [24]. Suhas B et al. , “Privacy sensitive speech analysis using federated learning to assess depression,” ICASSP, 2022.
- [25]. Yin Y et al. , “Speaker-invariant adversarial domain adaptation for emotion recognition,” in Proceedings of the 2020 International Conference on Multimodal Interaction, 2020, pp. 481–490.
- [26]. Li H et al., “Speaker-invariant affective representation learning via adversarial training,” in ICASSP. IEEE, 2020, pp. 7144–7148.
- [27]. Gat I et al. , “Speaker normalization for self-supervised speech emotion recognition,” arXiv preprint arXiv:2202.01252, 2022.
- [28]. Hsu W-N et al. , “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [29]. Busso C, Bulut M et al. , “Iemocap: Interactive emotional dyadic motion capture database,” *Language resources and evaluation*, vol. 42, no. 4, pp. 335–359, 2008.
- [30]. Busso C et al. , “Msp-improv: An acted corpus of dyadic interactions to study emotion perception,” *IEEE Transactions on Affective Computing*, vol. 8, no. 1, pp. 67–80, 2016.
- [31]. Baevski A et al. , “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 449–12 460, 2020.
- [32]. Ganin Y et al. , “Domain-adversarial training of neural networks,” *The journal of machine learning research*, vol. 17, no. 1, pp. 2096–2030, 2016.
- [33]. Valstar M et al. , “Avec 2016: Depression, mood, and emotion recognition workshop and challenge,” in Proceedings of the 6th international workshop on AVEC, 2016, pp. 3–10.
- [34]. Li Y 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. [PubMed: 22126712]
- [35]. Kroenke K 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. [PubMed: 18752852]
- [36]. Ter Smitten M et al., “Composite international diagnostic interview (CIDI), version 2.1,” Amsterdam: World Health Organization, no. 343, pp. 343–345, 1998.
- [37]. Bailey A et al., “Gender bias in depression detection using audio features,” in 2021 29th EUSIPCO. IEEE, 2021, pp. 596–600.
- [38]. Conneau A et al. , “Unsupervised cross-lingual representation learning for speech recognition,” arXiv preprint arXiv:2006.13979, 2020.
- [39]. von Platen P, “Wav2vec2.0 base,” 2022. [Online]. Available: <https://huggingface.co/facebook/wav2vec2-base-960h>
- [40]. Grosman J, “Xlsr wav2vec2 chinese (zh-cn),” <https://huggingface.co/jonatasgrosman/wav2vec2-large-xlsr-53-chinese-zh-cn>, 2021.
- [41]. Chinchor N, “Muc-4 evaluation metrics in proc. of the fourth message understanding conference 22–29,” 1992.
- [42]. Fukunaga K, *Introduction to statistical pattern recognition*. Elsevier, 2013.
- [43]. Guo J et al. , “Speaker verification using short utterances with dnn-based estimation of subglottal acoustic features.” in Interspeech, 2016, pp. 2219–2222.
- [44]. McNemar Q, “Note on the sampling error of the difference between correlated proportions or percentages,” *Psychometrika*, vol. 12, no. 2, pp. 153–157, 1947. [PubMed: 20254758]

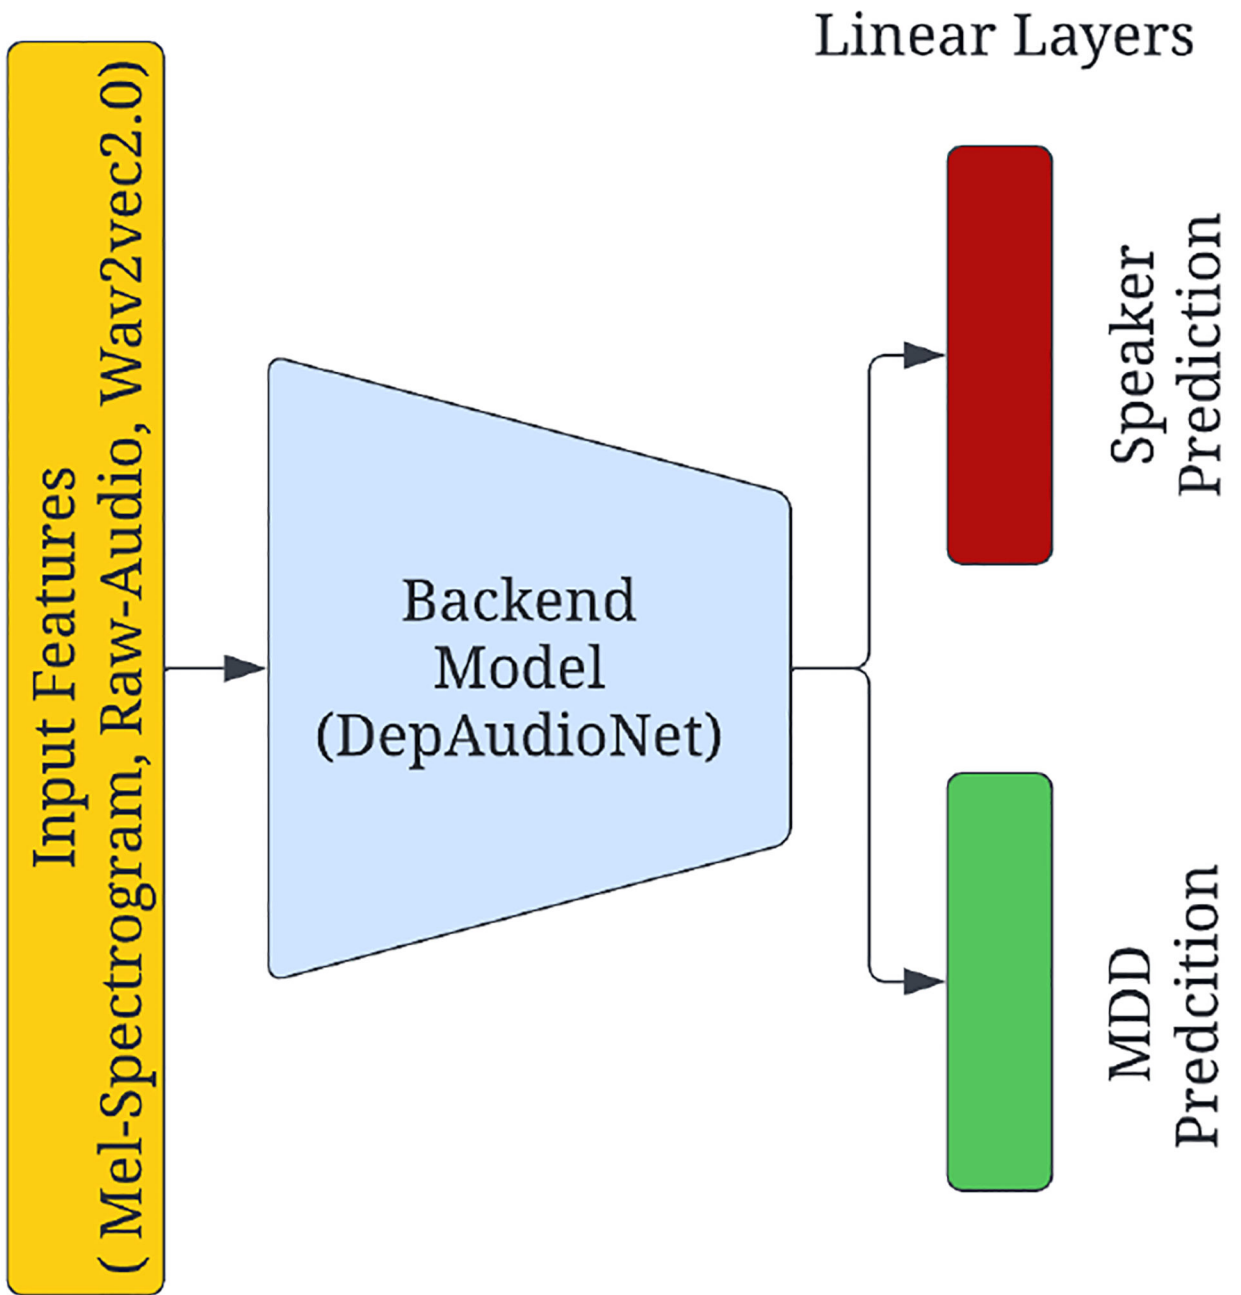


Figure 1:
Block diagram representing adversarial disentanglement of speaker and depression characteristics.

Table 1:

F1-scores for DepAudioNet using the DAIC-WOZ dataset for three features, with and without adversarial speaker disentanglement. F1-Avg is the average of F1-scores for non-depressed (F1-ND) and depressed (F1-D) classes. +Adv denotes adversarial training. λ values used for disentanglement are mentioned in parenthesis.

Input Feature	+Adv	F1-Avg	F1-ND	F1-D
Mel-spec	No	0.619	0.706	0.533
Mel-spec	Yes ($\lambda = 5e-6$)	0.646	0.732	0.560
Raw-audio	No	0.646	0.779	0.512
Raw-audio	Yes ($\lambda = 1e-6$)	0.660	0.726	0.594
Wav2vec2.0	No	0.686	0.804	0.567
Wav2vec2.0	Yes ($\lambda = 1e-3$)	0.692	0.808	0.576

Table 2:

F1-scores for DepAudioNet using the CONVERGE dataset for three features, with and without adversarial speaker disentanglement. ‘+Adv’ denotes adversarial training.

Input Feature	+Adv	F1-Avg	F1-ND	F1-D
Mel-spec	No	0.879	0.890	0.868
Mel-spec	Yes ($\lambda = 5e-5$)	0.890	0.903	0.877
Raw-audio	No	0.829	0.832	0.826
Raw-audio	Yes ($\lambda = 2e-4$)	0.857	0.870	0.844
XLSR-Mandarin	No	0.912	0.921	0.903
XLSR-Mandarin	Yes ($\lambda = 2e-4$)	0.915	0.925	0.906

Table 3:

Percentage change (relative) in F1-Avg after adversarial disentanglement for three input features using DAIC-WOZ and CONVERGE datasets.

Dataset	Input Feature	% Change in F1-Avg
DIAC-WOZ	Mel-spec	+4.36
	Raw-audio	+2.16
	Wav2vec2.0	+0.87
CONVERGE	Mel-spec	+1.25
	Raw-audio	+3.34
	XLSR-Mandarin	+0.33 *

* indicates that the change is not statistically significant.

Table 4:

J-ratios and F1-Avg scores for three input features with and without adversarial learning on DIAC-WOZ data. '+Adv' denotes adversarial training

Input Features	J-Ratio	F1-Avg
Mel-spec	4.81	0.619
Mel-spec+Adv	4.60	0.646
Raw-audio	2.17	0.646
Raw-audio+Adv	1.84	0.660
Wav2vec2.0	1.56	0.686
Wav2vec2.0+Adv	1.52	0.690

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 5:

J-ratios and F1-Avg scores for three input features with and without adversarial learning on CONVERGE data. '+Adv' denotes adversarial training

Input Features	J-Ratio	F1-Avg
Mel-spec	42.66	0.879
Mel-spec+Adv	40.61	0.890
Raw-audio	302.3	0.829
Raw-audio+Adv	74.64	0.857
XLSR-Mandarin	75.54	0.912
XLSR-Mandarin+Adv	71.62	0.915

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript