

Cost-effective Models for Detecting Depression from Speech

Mashrura Tasnim

Department of Computing Science, University of Alberta
Edmonton, AB, CA
mashrura@ualberta.ca

Jekaterina Novikova

Winterlight Labs
Toronto, ON, CA
jekaterina@winterlightlabs.com

Abstract—Depression is the most common psychological disorder and is considered as a leading cause of disability and suicide worldwide. An automated system capable of detecting signs of depression in human speech can contribute to ensuring timely and effective mental health care for individuals suffering from the disorder. Developing such automated system requires accurate machine learning models, capable of capturing signs of depression. However, state-of-the-art models based on deep acoustic representations require abundant data, meticulous selection of features, and rigorous training; the procedure involves enormous computational resources. In this work, we explore the effectiveness of two different acoustic feature groups - conventional hand-curated and deep representation features, for predicting the severity of depression from speech. We explore the relevance of possible contributing factors to the models' performance, including gender of the individual, severity of the disorder, content and length of speech. Our findings suggest that models trained on conventional acoustic features perform equally well or better than the ones trained on deep representation features at significantly lower computational cost, irrespective of other factors, e.g. content and length of speech, gender of the speaker and severity of the disorder. This makes such models a better fit for deployment where availability of computational resources is restricted, such as real time depression monitoring applications in smart devices.

Index Terms—Neural network, support vector machine, acoustic analysis, mental health

I. INTRODUCTION

Depression is a common psychological disorder. About 264 million people worldwide suffer from depression, which is almost 5% of the world's total population [1]. Only about 50% of the people experiencing major depression receive treatment. Due to lack of continuous monitoring and timely support, depression causes one death every 40 seconds, resulting in 800,000 deaths by suicide worldwide every year [1].

Conventional depression diagnostic systems, such as clinical assessment or standard questionnaires, require significant amount of time and active participation of the depressed individuals. Studies reveal that depression is reflected in behavioral fluctuations of certain day-to-day activities and physical parameters [2]. These findings have accelerated interventions in depression recognition using predictive models that incorporate input data of different modalities, among which audiovisual is one of the most explored areas. In this work, we emphasize on audio modality for its manifold benefits. Audio based depression detection system offers better privacy

for users of remote monitoring system. This kind of automated assessment takes only a few minutes of audio recording, therefore is less time-consuming, and would reduce burden on the individuals.

Multiple research efforts aim to develop a system that detects depression by analyzing the fluctuation of acoustic features in human speech ([3], [4]). An ML model that detects evidences of depression from audio data is a prerequisite for such a system. Existing best performing ML models that detect mental and cognitive diseases from audio data use either deep representation acoustic features, or a combination of conventional hand-crafted and deep features [5]. Although deep representation features offer a unified process of feature extraction, feature selection and model training, extracting and processing these features demands enormous computation resources including memory and processing time. This makes such models inconvenient for many real-world applications, where speed of data processing, model training and inference are of crucial importance [6]. Therefore, researchers and system designers need to make a choice of features when developing and deploying the model, considering both performance and cost. Some previous research compare the two approaches in the domain of cognitive disease detection [7] but to the best of our knowledge, no such research has been done so far in the domain of depression. To address this gap, in this work we have experimented with both hand-crafted conventional acoustic features and deep representation acoustic features. We address the following research questions:

- 1) *Between conventional and deep representation acoustic features, which ones are more effective in determining depression severity in terms of accuracy and computational cost?*
- 2) *Does the machine learning (ML) model performance vary based on gender of the subject?*
- 3) *What is the effect of content and length of speech data in predicting depression from speech?*

Answers to these questions enables the research community as well as system designers in making informed choice of modality, features, algorithm that suits best to the context, e.g. target user group and affordability. In this work, we compare performance of the ML models trained on each type of features extracted from speech samples of a variety of content and

length. Our key findings suggest that:

- 1) ML model trained on conventional acoustic feature set curated using expert domain knowledge demonstrates competitive performance as state-of-the-art models in predicting depression severity, irrespective of length and content of speech, and gender of the speaker
- 2) Usage of deep representation features resulted in marginal improvement of performance (0.0004%) consuming 1000 times more memory and 3000 times more computation time.

As such, we claim that models predicting depression from human speech that are trained on conventional acoustic features are a better choice than the models trained on deep acoustic representations in the situations when computational resources are limited, e.g. in mental healthcare applications for portable or wearable devices. On the other hand, deep representation models fit better to the scenarios where abundant training data is available, for example social media, and computational resources is a legitimate trade-off for better performance.

II. RELATED WORKS

Individuals suffering from psychological and neurological disorders like depression exhibit measurable fluctuation in vocal parameters ([8] and [9]). Significant number of research have been conducted to relate these parameters with depression severity. DAIC-WoZ dataset [10] is a widely used dataset in acoustic based depression severity prediction, consisting of structured interviews of participants conducted by a virtual agent. Two subsets of this dataset have been introduced as the challenge corpus of three Audio/Visual Emotion Challenges (AVEC) in 2016 [11], 2017 [12] and 2019 [13], where participants proposed machine learning models to predict depression score on the PHQ-8 scale [14]. Handcrafted acoustic features have been exploited for this task for the last few decades, while deep representation of acoustic features have become popular in recent years. Further, we present a summary of existing works in this area and compare them based on the type of acoustic features.

A. Conventional Acoustic Features

Conventional acoustic features fall in temporal, spectral, energy and voicing related categories, from which researchers hand-pick the ones that are most suitable for predicting certain disorders, such as depression [9]. Over the time, certain sets of these acoustic features, introduced in speech emotion and depression recognition challenges, have gained popularity, among which baseline feature sets of AVEC 2013 [15] and AVEC 2016 [11], INTERSPEECH ComParE [16], extended Geneva Minimalistic Acoustic Parameter Set (eGeMAPS) [17] are noteworthy. Development of feature extraction toolkits like openSMILE [18], COVAREP [19] has made it easier for researchers to extract these features for the purpose of speech analysis in different aspects.

B. Deep Representation Features

Deep representation of acoustic features are inspired by the deep learning paradigms common in image processing. Here, spectral images of speech instances are fed into pre-trained image recognition CNNs and a set of the resulting activations are extracted as feature vectors. In AVEC 2019 Depression Detection Sub-challenge (DSC), Deep representation features from four robust pre-trained CNNs using VGG-16 [20], AlexNet [21], DenseNet-121, and DenseNet-201 [22] were included as challenge baseline features. Participants chose between using one or more sets of deep representation features ([23], [24]), and combining them with traditional features [5] and obtained competitive performances (Table III). Deep representation provides the option to unite feature extraction, feature selection and model training into a single automated generalisable procedure, compromising the opportunity to incorporate expert domain knowledge, and necessitating considerably higher computational cost.

C. Depression Detection Models

AVEC 2016 challenge dataset was used in analysis presented in [25], [26], [27], [28], [29], [30], [31], [32] and [33]. Williamson *et al.* extracted formants, MFCCs, glottal features, loudness [25]. In addition to COVAREP [19] audio features, [26] took text topics into account, while [27] considered a more extended set of features of audio, video and text modalities. [29] added Delta and Delta-Delta coefficients, mean, median, standard deviation, peak-magnitude to RMS ratio to the set of challenge baseline audio features. Applying similar higher-order statistics on the baseline features, [30] constructed an extended feature set of 553 features, of which they identified 279 features with statistically significant univariate correlation. [31] implemented multi-modal sentence-level embedding on log-Mel spectrogram and MFCC features. In their recent work, [34] exploited a combination of eGeMAPS and INTERSPEECH features extracted from the longest ten segments of each audio sample. They reshape the feature vector in an image-like 2D feature map in row-major order and adopted Deep Convolutional Generative Adversarial Net (DCGAN) for feature vector generation. [33] trained three spectrogram-based Deep Neural Network architectures phoneme consonant and vowel units and their fusion. Their findings suggest that deep learned consonant-based acoustic characteristics lead to better recognition results than vowel-based ones, and the fusion of vowel and consonant speech characteristics outperforms the other models on the task. [32] described a transfer attention mechanisms from speech recognition to aid depression severity measurement. The transfer is applied in a two-level hierarchical network which reflecting the natural hierarchical structure of speech.

The AVEC 2016 challenge corpus included training, development and test partitions of audio samples. Using acoustic features exclusively, the lowest root-mean-square-error (RMSE) of 5.52 and 6.42 on the development and test set were reported in [35] and [36] respectively. RMSE 4.99, 5.40, 5.66 and 6.42 were reported in [27], [35], [32] and [36] respectively

on the test set. The challenge baseline RMSE 6.74 (mean absolute error (MAE) 5.36) and 7.78 (MAE 5.72) were set for the development and test set respectively [12].

Comparatively, fewer number of studies have been conducted on the recently published AVEC 2019 DDS dataset; which is a super-set of the AVEC 2016 dataset. A wide range of audio features has been provided as the challenge baseline encompassing both handcrafted sets of conventional features and deep representation of acoustic features, including eGeMAPS, Mel Frequency Cepstral Coefficients (MFCCs), Bag of Audio Words (BoAW) and two sets of deep spectrum features by feeding spectral images of speech instances into pre-trained image recognition Convolutional Neural Networks (CNN) (VGG-16 [22] and DenseNet-121 [20]) and extracting the resulting activations as feature vectors. Of these, deep spectrum features were used by [24] and [23], and MFCCs and eGeMAPS by [37]. [5] exploited all the baseline feature sets, while [38] extracted the AVEC 2017 baseline feature set [12] using the COVAREP software toolbox [19], in addition to eGeMAPS. Different configurations of long short-term memory (LSTM) networks were used in all of these works except Zhang *et al.* [38], who adopted random forest and logistic regression. Acoustic models proposed by [5] and [38] achieved lowest RMSE of 5.11 and 6.78 on the development and test partitions, respectively.

III. MATERIALS AND METHODS

A. Dataset

We use DEPressure and Anxiety Crowdsourced corpus (DEPAC) [39] in this experiment. The dataset consists of 2,674 audio samples collected from 571 subjects located in Canada and the United States. 54.67% of the study subjects are female and 45.33% are male, aged between 18 and 76 years, and they received 1 to 26 years of formal education. The data was collected via crowdsourcing and consists of a variety of self-administered speech tasks (Table I). The participants completed these tasks using Amazon Mechanical Turk (mTurk) ¹. The speech tasks were curated to increase phonemic variety and were supported by literature on detecting mental disorders, such as Alzheimer’s Disease (AD) [40] and depression [41], [42], [43] from speech.

In this dataset, the depression severity is represented by Patient Health Questionnaire (PHQ-9) scores. It is a 3-point self-rated measure for depressive symptoms, including 9 questions. To ensure comparability of our results with works done on popular subsets of DAIC-WoZ corpus [10], i.e. AVEC 2017 [12] and AVEC 2019 [13], we used responses to 8 PHQ questions in our analysis and reported our results on PHQ-8 scores. The score ranges from 0 to 24 on PHQ-8 scale where a score higher than 5, 9 and 14 represent mild, moderate and severe level of depression respectively. The mean PHQ-8 score of DEPAC corpus (M) is 6.56 with standard deviation (SD) of 5.56.

¹<https://www.mturk.com>

TABLE I: Speech tasks in DEPAC corpus

Speech Task	Description	Average Duration
Phoneme Task	Record “aah” sound for as long as the participant could hold breath	5.79 sec
Phonemic Fluency	Pronounce as many unique words as possible starting with the letters “F”, “A” or “S”	22.13 sec
Picture Description	Describe a picture shown on the screen	46.60 sec
Semantic Fluency	Describe a positive experience they expected to have within five years in future	43.76 sec
Prompted Narrative	Tell a personal story, describing the day, a hobby, or a travel experience	45.34 sec

B. Audio Quality Enhancement

To suppress possible background noise present in the samples and improve quality of the audio, we applied *logmmse* enhancement technique [44] on the audio samples. This method was found the best among existing audio enhancement algorithms [45]. The enhancement step is found statistically significant ($p \leq 0.005$) on 94% of the 220 conventional acoustic features in Wilcoxon signed-rank test with Bonferroni correction.

Audio volume was normalized to -20 dBFS across all speech segments to control for variation caused by recording conditions such as microphone placement.

C. Acoustic Features

We extracted two sets of acoustic features, representing hand-crafted sets of conventional features and deep learning features:

1) *Conventional acoustic features*: This set included 220 acoustic features, extracted from each audio sample. The feature set includes:

- **Spectral features**: Intensity (auditory model based), MFCC 0-12, Zero-Crossing Rate (ZCR)
- **Voicing-related features**: Fundamental frequency (F_0), Harmonic-to-Noise Ratio (HNR), shimmer and jitter, durational features, pauses and fillers, phonation rate

Statistical functionals including minimum, maximum, average, and variance were computed on the low-level descriptors. Additionally, skewness and kurtosis were calculated on MFCCs, first and second order derivatives of MFCCs, and Zero Crossing Rate (ZCR) [46].

A Python implementation of Praat phonetic analysis toolkit [47] has been used to extract the majority of these features. The MFCC features and their functionals were computed using `python_speech_features`² library.

2) *Deep Representation Features*: Deep representation of acoustic features are inspired by the deep learning paradigms common in image processing. Here, spectral images of speech instances are fed into pre-trained image recognition CNNs and a set of the resulting activations are extracted as feature

²https://pypi.org/project/python_speech_features/

vectors. VGG-16 is a type of Convolutional Neural Network (CNN) which is considered to be one of the best computer vision models to date. We used DeepSpectrum library [48] to extract features from a pre-trained VGG-16 CNN [22]. The speech files are first transformed into mel-spectrogram images with 128 mel-frequency bands. Then, the spectral images are forwarded through the pre-trained networks. A 4,096-dimensional feature vector is then formed from the activations of the second fully connected layer in VGG-16. The features were extracted at a window width of 1s and a hop size of 300 ms from each audio sample.

D. Data Preprocessing

1) *Standardization* : The range of values of audio features tends to vary widely. To ensure even contribution of all features in the regression task, and to speed up gradient descent convergence of the deep neural network, once acoustic features were extracted from the audio samples, we standardized them using z-scores, i.e., subtracting the mean and dividing by standard deviation. The standard score of a sample x of feature f_i is calculated as:

$$y = \frac{x - \mu}{\sigma} \quad (1)$$

here μ and σ are the mean and standard deviation of the values of f_i in all training samples.

2) *Feature Selection*: We applied minimum Redundancy-Maximum Relevance (mRMR) algorithm [49] to select the most relevant features with respect to the PHQ scores, minimizing redundancy in the selected set of features. 10% features were selected from each set for training the ML models.

E. Model Training

Following [7] and [4], we train a combination of linear and non-linear ML models separately on conventional and deep learning acoustic features:

- Support Vector Machines (SVM): Radial Basis Function (RBF) kernel SVM was trained. Values of hyperparameters ‘C’ and ‘gamma’ were tuned by 5-fold grid-search cross validation (cv).
- Random Forest (RF): Scikit Learn implementation of random forest regressor was used. Number of estimator trees and maximum depth were tuned through grid-search cv.
- Feedforward Neural Network (FNN): The FNN model consists of 4 hidden layers, with 500 hidden units on the first layer, 250 in the second and 125 in the rest of the hidden layers. 30% dropout on the output of each of the hidden layers of the FNN. We use Adam optimizer in all FNN models with the learning rate of 0.001. Each of the FNN models is trained for 150 epochs.

The discussion presented by [7] suggest that for small audio corpus like the ADReSSo challenge dataset (237 samples) [50], either leave-one-subject-out cross validation or k-fold cross validation can be applied. However, the dataset used in this work is considerably larger than the ADReSSo challenge dataset. Considering the size of the dataset and corresponding

computational complexity, we decided to report evaluation metrics with 5-fold cross-validation (CV) for the models. We create 5 subject-independent folds, train the model using 4 of them, and use the rest for testing. We repeat the process for all 5 folds and report evaluation metrics averaging across predictions on all the folds. These folds preserve the same ratio of depression severity in each training and test partitions.

To understand the effect of speech content on ML models’ performance, we separated samples with each type of speech task and trained models on each type of them. We repeated the same process for conventional and VGG-16 features.

We ran our experiment on a MacBook Pro with Intel Core i7 CPU at clock speed of 2.67 GHz. The system availed 16 GB memory. The data preprocessing and model training was done in Python programming language.

TABLE II: Regression error of models trained on conventional and VGG-16 features

Algorithm	Gender	RMSE		MAE	
		Conventional	VGG-16	Conventional	VGG-16
SVM	Male	5.04	7.89	4.22	6.95
	Female	5.64	7.11	4.33	6.23
	Overall	5.38	7.48	4.28	6.56
RF	Male	5.15	5.06	4.37	4.27
	Female	5.47	5.51	4.34	4.32
	Overall	5.32	5.31	4.31	4.33
FNN	Male	5.10	5.19	4.45	4.30
	Female	5.54	5.67	4.51	4.34
	Overall	5.35	5.46	4.40	4.32

IV. RESULT AND DISCUSSION

Here we present the performance of the ML models trained on DEPAC dataset with a view to find answers to the research questions (RQs), outlined in Section I.

A. Effectiveness of different types of acoustic features in measuring depression

We trained 3 different ML models on each type of acoustic feature, i.e. conventional and VGG-16. We report the RMSE and MAE error of each model trained separately on samples from male and female subjects, along with the overall performance on the entire dataset (Table II). We compare the performance of our best model with the state-of-the-art (Table III). We report CPU time required to train each model to assist future researchers and system designers in making informed choice of feature type and ML model.

1) *Performance of models trained on conventional and deep representation features*: SVM and FNN models performed better on conventional features than on VGG-16, while performance of RF is marginally better (0.0004%) on VGG-16 (Table II). These findings are consistent with the previous works presented in Table III. Conventional features presented in [34] modeled depression marginally better than models with deep representation features [23], [32].

In comparison to the state of the art acoustic models, our proposed RF models show competitive performance. The RF

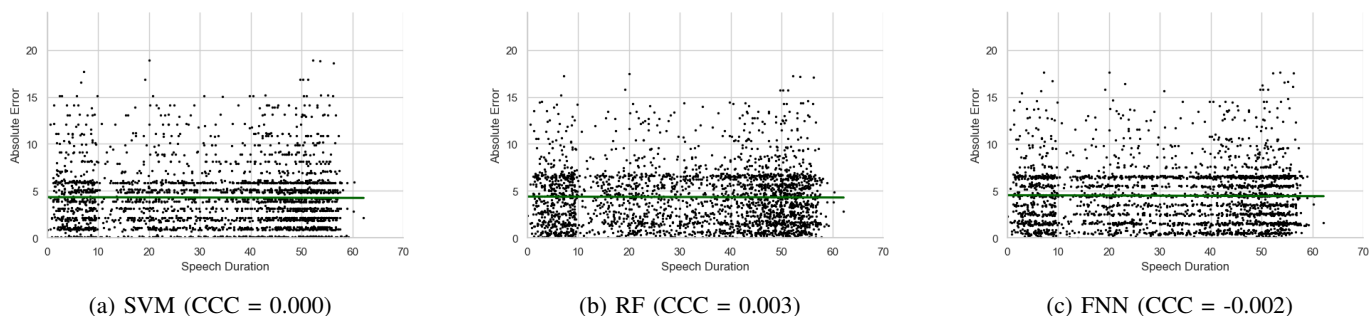


Fig. 1: Correlation between speech length and prediction error of models trained on conventional acoustic features

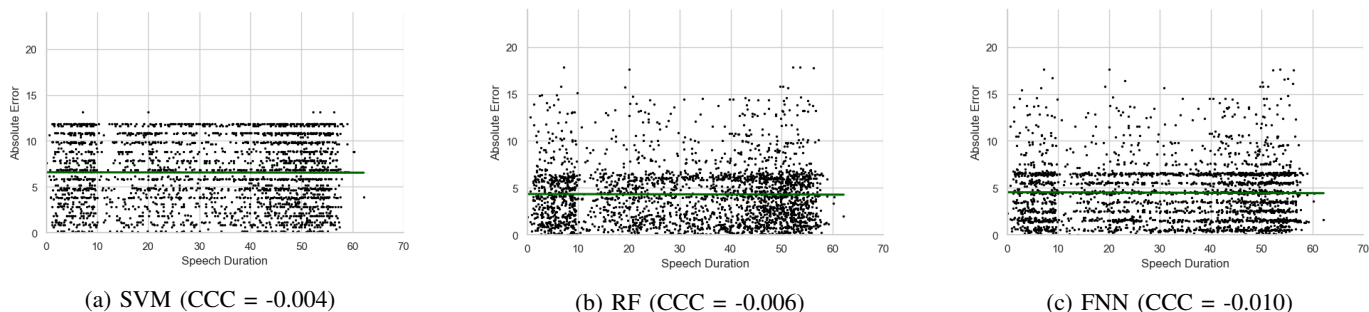


Fig. 2: Correlation between speech length and prediction error of models trained on VGG-16 features

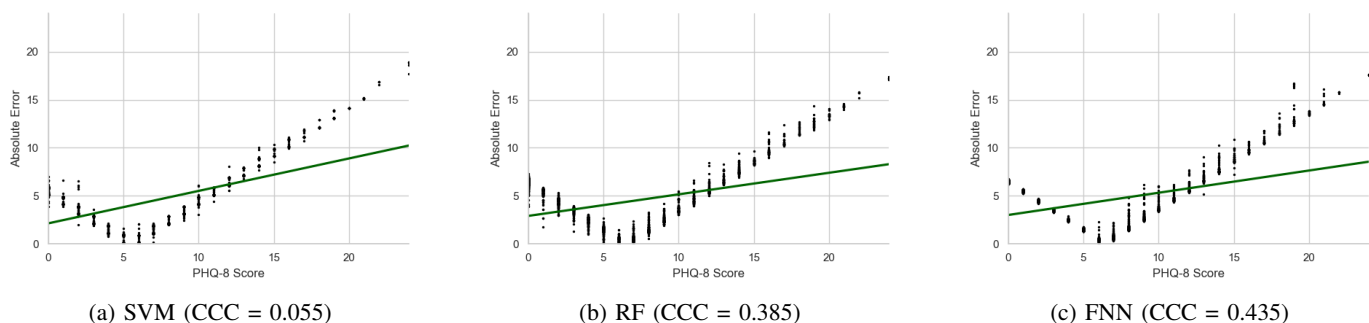


Fig. 3: Correlation between depression severity and prediction error of models trained on conventional acoustic features

model trained on both types of features outperforms almost all the existing works reporting similar performance metrics on PHQ-8 scale (see Table III). Only [5] reported lower RMSE than us, fusing all four sets of AVEC 2019 baseline features, which is a combination of conventional (MFCC, Bag-of-Audio Words, eGeMAPS) and deep representation (VGG-16) features, and formulating a multi-level LSTM architecture. Our proposed RF model trained on conventional features produces competitive performance to their proposed model, while substantially decreasing computational requirements. The VGG-16 features collected in the same manner as described by [5] from our audio corpus occupy 11.21 GB of memory, while our presented conventional feature file size is only 11 MB. Preprocessing and training models on conventional acoustic features took on average 3 minutes, while the procedure on VGG-16 features took at least 150 hours on the same computational environment (2.6 GHz 6 core Intel Core i7 processor,

16 GB memory). In short, our RF model using VGG-16 features offers 0.0004% improvement in performance than the same model using conventional features, using 1000 times more memory and 3000 times more processing time, implying similar or more computational resource is required for training complex models on multimodal features for marginal performance improvement. Therefore, the conventional features provided better opportunity to adjust model parameters for performance improvement.

Results (Table III) demonstrate that, in most cases, RMSE and MAE are lower for male subjects than female subjects. The reason behind this can be the lower severity of depression among male subjects than females in DEPAC dataset [39]. The skewness in the dataset causes bias in model prediction, as described in [52]. For real world applications, this issue needs to be taken care of by ensuring gender balance in training data.

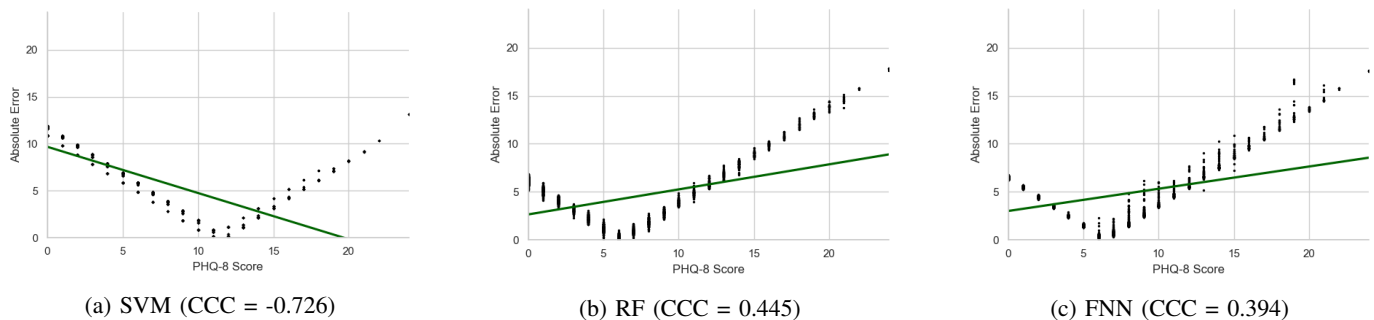


Fig. 4: Correlation between depression severity and prediction error of models trained on VGG-16 features

TABLE III: Comparison of performance of SOTA ML models trained on different combinations of features. Bold denotes regression error of our proposed model.

Feature Type	Study	RMSE	MAE
Conventional	Formants, MFCCs, glottal features, loudness, AVEC 2017 dataset [51]	6.38	5.32
	COVAREP feature set, AVEC 2017 dataset [36]	6.34	5.30
	COVAREP features and functional, AVEC 2016 dataset)	6.50	5.13 [30]
	MFCC, AVEC 2016 dataset [31]	5.78	-
	MFCC and eGeMAPS features, AVEC 2019 dataset [37]	6.20	-
	eGeMAPS, INTERSPEECH features, AVEC 2016 dataset [34]	5.52	4.63
	eGeMAPS and COVAREP features, AVEC 2019 dataset [38]	6.78	5.77
Deep representation	VGG-16 features, AVEC 2019 dataset [23]	5.70	-
	Mel-spectra, AVEC 2017 dataset [32]	5.66	4.28
Conventional + deep combined	MFCC, BoAW, eGeMAPS and VGG-16 features, AVEC 2019 dataset [5]	5.11	-
Conventional	MFCCs, HNR, jitter, shimmer, ZCR features, DEPAC dataset [39]	5.31	4.33

2) *Significance of performance deviation of models trained on conventional and deep representation features:* We performed two-sample t-tests to identify if the performance deviations of the models are significant when trained on conventional acoustic features and VGG-16 features.

There was a significant difference between absolute errors in predictions of SVM models trained on WLL acoustic features ($M = 4.28, SD = 3.25$) and VGG-16 features ($M = 6.5, SD = 3.59$); $t(5332) = -24.20, p = 6.86e - 123 < .05$. The absolute errors in predictions of SVM with model is significantly higher when trained on VGG-16 features than when trained on conventional features.

On the other hand, there was no significant difference

TABLE IV: Time elapsed in different stages of model training

Processing step	Algorithm	Conventional	VGG-16
Data loading	-	1.132	244450
Preprocessing	-	96.834	545483
Model training	SVM	1.600	5593
	RF	0.715	981
	FNN	220.853	10270
Prediction	SVM	0.412	53
	RF	0.040	9
	FNN	1.102	7
Total	SVM	99.978	795579
	RF	98.715	790923
	FNN	267.227	800210

between absolute errors in predictions of our best performing RF and FNN models trained on conventional acoustic features ($RF : M = 4.34, SD = 3.09$; $FNN : M = 4.32, SD = 3.14$) and VGG-16 features ($RF : M = 4.31, SD = 3.11$; $FNN : M = 4.48, SD = 3.11$). The test scores of RF ($t(5332) = 0.38, p = .70 > .05$) and FNN ($t(5332) = -1.81, p = .07 > .05$) indicate that the deviation of errors in prediction of the models are not significantly different irrespective of training features.

B. Effect of speech task type on ML model performance

The results (Table V) do not reflect any significant deviation of model performance on the basis of speech task, therefore it is possible to recommend as a design choice any speech task of a similar length and content.

C. Correlation between model performance and speech length

From Figure 1 and 2 one can see that no significant correlation is observed between model performance (absolute error of each prediction) and length of corresponding sample. The Concordance Correlation Coefficient (CCC) score for SVM, RF and FNN models are 0.000, 0.003 and -0.002 for conventional features and -0.004, -0.006 and -0.010 for VGG-16 features respectively. The near-zero CCC values indicate that in the case of our dataset, the speech length of samples does not influence the models' performance. Note that all speech samples in DEPAC are less than one minute.

TABLE V: Regression error of models trained on speech samples of different tasks

Algorithm	Speech task	RMSE		MAE	
		Conventional	VGG-16	Conventional	VGG-16
SVM	Semantic fluency	5.30	6.62	4.21	5.77
	Prompted narrative	5.29	6.63	4.24	5.77
	Phoneme task	5.49	6.40	4.33	5.54
	Phonemic fluency	5.45	6.49	4.35	5.62
	Picture description	5.43	6.54	4.36	5.67
RF	Semantic fluency	5.24	5.24	4.25	4.24
	Prompted narrative	5.31	5.25	4.30	4.25
	Phoneme task	5.39	5.29	4.37	4.30
	Phonemic fluency	5.38	5.29	4.38	4.30
	Picture description	5.42	5.31	4.40	4.31
FNN	Semantic fluency	5.34	7.13	4.28	5.34
	Prompted narrative	5.30	7.13	4.30	5.35
	Phoneme task	5.49	7.33	4.39	5.50
	Phonemic fluency	5.39	7.30	4.33	5.47
	Picture description	5.45	7.29	4.45	5.47

D. Correlation between depression severity and model performance

Absolute errors for each sample are plotted against ground truth PHQ-8 score for the models trained on conventional and VGG-16 features in Figure 3 and 4. The CCC score for SVM, RF and FNN models are 0.055, 0.385 and 0.435 for conventional features and -0.726, 0.445 and 0.394 for VGG-16 features respectively. The plots, along with high positive CCC scores for most of the models, imply that the samples with higher PHQ-8 scores contribute more to the overall prediction error of the models. This is caused by the imbalance in the number of samples with high and low PHQ-8 scores in DEPAC dataset [39]. The higher density of samples with subthreshold (≤ 5) PHQ-8 score bias the models to make predictions close to the mean PHQ-8 (6.56) of the dataset. This observation strengthens the necessity of balancing the samples in training models to be used in real world application.

V. CONCLUSION AND FUTURE WORKS

Speech has proven to be a reliable marker for depression assessment. But in order to deploy a machine learning model in a practical system, it is necessary to identify the most informative acoustic feature, along with an efficient and cost-effective process to train the model. In this paper, we study

the performance of conventional acoustic feature-based and pre-trained deep representation based models on predicting depression severity from speech. We observe that the hand-curated feature based approach achieves better performance in terms of lower RMSE and MAE, at a remarkably less computation time. Our experiments show that gender of the speaker and distribution of score affect the model performance, and should be taken care of while formulating balanced training data. We also report that content and length of speech do not show significant impact as long as the length of speech samples is reasonably short, less than one minute in our case. To summarize, we suggest using ML models trained on conventional features in resource-limited real-time situations and deep models in scenarios where fine-grained analysis involving higher computational power is crucial. In our future work, we plan to explore generalizability of the findings across other datasets and disorders.

REFERENCES

- [1] "Depression," 2020, <https://www.who.int/news-room/fact-sheets/detail/depression>.
- [2] R. Wang, F. Chen, Z. Chen, T. Li, G. Harari, S. Tignor, X. Zhou, D. Ben-Zeev, and A. T. Campbell, "Studentlife: assessing mental health, academic performance and behavioral trends of college students using smartphones," in *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 2014, pp. 3–14.
- [3] X. Ma, H. Yang, Q. Chen, D. Huang, and Y. Wang, "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.
- [4] M. Tasnim and E. Stroulia, "Detecting depression from voice," in *Canadian Conference on Artificial Intelligence*. Springer, 2019, pp. 472–478.
- [5] A. Ray, S. Kumar, R. Reddy, P. Mukherjee, and R. Garg, "Multi-level attention network using text, audio and video for depression prediction," in *Proceedings of the 9th international on audio/visual emotion challenge and workshop*, 2019, pp. 81–88.
- [6] B. Yalamanchili, N. S. Kota, M. S. Abbaraju, V. S. S. Nadella, and S. V. Alluri, "Real-time acoustic based depression detection using machine learning techniques," in *2020 International conference on emerging trends in information technology and engineering (ic-ETITE)*. IEEE, 2020, pp. 1–6.
- [7] A. Balagopalan and J. Novikova, "Comparing Acoustic-Based Approaches for Alzheimer's Disease Detection," in *Proc. Interspeech 2021*, 2021, pp. 3800–3804.
- [8] J. K. Darby, N. Simmons, and P. A. Berger, "Speech and voice parameters of depression: A pilot study," *Journal of Communication Disorders*, vol. 17, no. 2, pp. 75–85, 1984.
- [9] N. Cummins, V. Sethu, J. Epps, and J. Krajewski, "Probabilistic acoustic volume analysis for speech affected by depression," in *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [10] J. Gratch, R. Artstein, G. Lucas, G. Stratou, S. Scherer, A. Nazarian, R. Wood, J. Boberg, D. DeVault, S. Marsella *et al.*, "The distress analysis interview corpus of human and computer interviews," in *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, 2014, pp. 3123–3128.
- [11] M. Valstar, J. Gratch, B. Schuller, F. Ringeval, D. Lalanne, M. Torres Torres, S. Scherer, G. Stratou, R. Cowie, and M. Pantic, "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.
- [12] F. Ringeval, B. Schuller, M. Valstar, J. Gratch, R. Cowie, S. Scherer, S. Mozgai, N. Cummins, M. Schmitt, and M. Pantic, "Avec 2017: Real-life depression, and affect recognition workshop and challenge," in *Proceedings of the 7th Annual Workshop on Audio/Visual Emotion Challenge*, 2017, pp. 3–9.

- [13] F. Ringeval, B. Schuller, M. Valstar, N. Cummins, R. Cowie, L. Tavabi, M. Schmitt, S. Alisamir, S. Amiriparian, E.-M. Messner *et al.*, “Avec 2019 workshop and challenge: state-of-mind, detecting depression with ai, and cross-cultural affect recognition,” in *Proceedings of the 9th International on Audio/Visual Emotion Challenge and Workshop*, 2019, pp. 3–12.
- [14] K. Kroenke, R. L. Spitzer, and J. B. Williams, “The phq-9: validity of a brief depression severity measure,” *Journal of general internal medicine*, vol. 16, no. 9, pp. 606–613, 2001.
- [15] M. Valstar, B. Schuller, K. Smith, F. Eyben, B. Jiang, S. Bilakhia, S. Schnieder, R. Cowie, and M. Pantic, “Avec 2013: the continuous audio/visual emotion and depression recognition challenge,” in *Proceedings of the 3rd ACM international workshop on Audio/visual emotion challenge*, 2013, pp. 3–10.
- [16] B. Schuller, S. Steidl, A. Batliner, A. Vinciarelli, K. Scherer, F. Ringeval, M. Chetouani, F. Wenginger, F. Eyben, E. Marchi *et al.*, “The interspeech 2013 computational paralinguistics challenge: Social signals, conflict, emotion, autism,” in *Proceedings INTERSPEECH 2013, 14th Annual Conference of the International Speech Communication Association, Lyon, France*, 2013.
- [17] F. Eyben, *Real-time speech and music classification by large audio feature space extraction*. Springer, 2015.
- [18] F. Eyben, M. Wöllmer, and B. Schuller, “Opensmile: the munich versatile and fast open-source audio feature extractor,” in *Proceedings of the 18th ACM International Conference on Multimedia*. ACM, 2010, pp. 1459–1462.
- [19] G. Degottex, J. Kane, T. Drugman, T. Raitio, and S. Scherer, “Covarep—a collaborative voice analysis repository for speech technologies,” in *2014 IEEE International Conference on acoustics, speech and signal processing (icassp)*. IEEE, 2014, pp. 960–964.
- [20] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” in *Proceedings of the IEEE Conference on computer vision and pattern recognition*, 2017, pp. 4700–4708.
- [21] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” *Advances in neural information processing systems*, vol. 25, pp. 1097–1105, 2012.
- [22] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [23] M. Rodrigues Makiuchi, T. Warnita, K. Uto, and K. Shinoda, “Multi-modal fusion of bert-cnn and gated cnn representations for depression detection,” in *Proceedings of the 9th International on Audio/Visual Emotion Challenge and Workshop*, 2019, pp. 55–63.
- [24] S. Yin, C. Liang, H. Ding, and S. Wang, “A multi-modal hierarchical recurrent neural network for depression detection,” in *Proceedings of the 9th International on Audio/Visual Emotion Challenge and Workshop*, 2019, pp. 65–71.
- [25] J. R. Williamson, E. Godoy, M. Cha, A. Schwarzentruher, P. Khorrani, Y. Gwon, H.-T. Kung, C. Dagli, and T. F. Quatieri, “Detecting depression using vocal, facial and semantic communication cues,” in *Proceedings of the 6th International Workshop on Audio/Visual Emotion Challenge*, 2016, pp. 11–18.
- [26] B. Sun, Y. Zhang, J. He, L. Yu, Q. Xu, D. Li, and Z. Wang, “A random forest regression method with selected-text feature for depression assessment,” in *Proceedings of the 7th Annual Workshop on Audio/Visual Emotion Challenge*. ACM, 2017, pp. 61–68.
- [27] Y. Gong and C. Poellabauer, “Topic modeling based multi-modal depression detection,” in *Proceedings of the 7th Annual Workshop on Audio/Visual Emotion Challenge*. ACM, 2017, pp. 69–76.
- [28] L. Yang, H. Sahli, X. Xia, E. Pei, M. C. Oveneke, and D. Jiang, “Hybrid depression classification and estimation from audio video and text information,” in *Proceedings of the 7th Annual Workshop on Audio/Visual Emotion Challenge*. ACM, 2017, pp. 45–51.
- [29] A. Samareh, Y. Jin, Z. Wang, X. Chang, and S. Huang, “Predicting depression severity by multi-modal feature engineering and fusion,” in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [30] T. Al Hanai, M. M. Ghassemi, and J. R. Glass, “Detecting depression with audio/text sequence modeling of interviews,” in *Interspeech*, 2018, pp. 1716–1720.
- [31] A. Haque, M. Guo, A. S. Miner, and L. Fei-Fei, “Measuring depression symptom severity from spoken language and 3d facial expressions,” *arXiv preprint arXiv:1811.08592*, 2018.
- [32] Z. Zhao, Z. Bao, Z. Zhang, N. Cummins, H. Wang, and B. Schuller, “Hierarchical attention transfer networks for depression assessment from speech,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7159–7163.
- [33] M. Muzammel, H. Salam, Y. Hoffmann, M. Chetouani, and A. Othmani, “Audvowelconsnet: A phoneme-level based deep cnn architecture for clinical depression diagnosis,” *Machine Learning with Applications*, vol. 2, p. 100005, 2020.
- [34] L. Yang, D. Jiang, and H. Sahli, “Feature augmenting networks for improving depression severity estimation from speech signals,” *IEEE Access*, vol. 8, pp. 24 033–24 045, 2020.
- [35] L. Yang, D. Jiang, W. Han, and H. Sahli, “Dcnn and dnn based multi-modal depression recognition,” in *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 2017, pp. 484–489.
- [36] Z. S. Syed, K. Sidorov, and D. Marshall, “Depression severity prediction based on biomarkers of psychomotor retardation,” in *Proceedings of the 7th Annual Workshop on Audio/Visual Emotion Challenge*, 2017, pp. 37–43.
- [37] W. Fan, Z. He, X. Xing, B. Cai, and W. Lu, “Multi-modality depression detection via multi-scale temporal dilated cnns,” in *Proceedings of the 9th International on Audio/Visual Emotion Challenge and Workshop*, 2019, pp. 73–80.
- [38] L. Zhang, R. Duvvuri, K. K. Chandra, T. Nguyen, and R. H. Ghomi, “Automated voice biomarkers for depression symptoms using an online cross-sectional data collection initiative,” *Depression and anxiety*, 2020.
- [39] M. Tasnim, M. Eghghaghi, B. Diep, and J. Novikova, “Depac: a corpus for depression and anxiety detection from speech,” in *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology (CLPsych) (in press)*. Association for Computational Linguistics, 2022.
- [40] J. G. Borkowski, A. L. Benton, and O. Spreen, “Word fluency and brain damage,” *Neuropsychologia*, vol. 5, no. 2, pp. 135–140, 1967.
- [41] H. Jiang, B. Hu, Z. Liu, L. Yan, T. Wang, F. Liu, H. Kang, and X. Li, “Investigation of different speech types and emotions for detecting depression using different classifiers,” *Speech Communication*, vol. 90, pp. 39–46, 2017.
- [42] P. Fossati, A.-M. Ergis, J.-F. Allilaire *et al.*, “Qualitative analysis of verbal fluency in depression,” *Psychiatry research*, vol. 117, no. 1, pp. 17–24, 2003.
- [43] R. N. Trifu, B. NEMEŞ, C. Bodea-Hătegan, and D. Cozman, “Linguistic indicators of language in major depressive disorder (mdd). an evidence based research,” *Journal of Evidence-Based Psychotherapies*, vol. 17, no. 1, 2017.
- [44] Y. Ephraim and D. Malah, “Speech enhancement using a minimum mean-square error log-spectral amplitude estimator,” *IEEE transactions on acoustics, speech, and signal processing*, vol. 33, no. 2, pp. 443–445, 1985.
- [45] Y. Hu and P. C. Loizou, “Subjective comparison of speech enhancement algorithms,” in *2006 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*, vol. 1. IEEE, 2006, pp. 1–1.
- [46] D. M. Low, K. H. Bentley, and S. S. Ghosh, “Automated assessment of psychiatric disorders using speech: A systematic review,” *Laryngoscope Investigative Otolaryngology*, vol. 5, no. 1, pp. 96–116, 2020.
- [47] P. Boersma and V. Van Heuven, “Speak and unspeak with praat,” *Glott International*, vol. 5, no. 9/10, pp. 341–347, 2001.
- [48] S. Amiriparian, M. Gerczuk, S. Ottl, N. Cummins, M. Freitag, S. Pugachevskiy, A. Baird, and B. Schuller, “Snore sound classification using image-based deep spectrum features,” in *Interspeech 2017*. ISCA, Aug. 2017, pp. 3512–3516.
- [49] Z. Zhao, R. Anand, and M. Wang, “Maximum relevance and minimum redundancy feature selection methods for a marketing machine learning platform,” in *2019 IEEE international conference on data science and advanced analytics (DSAA)*. IEEE, 2019, pp. 442–452.
- [50] S. Luz, F. Haider, S. de la Fuente, D. Fromm, and B. MacWhinney, “Detecting cognitive decline using speech only: The addresso challenge,” *arXiv preprint arXiv:2104.09356*, 2021.
- [51] J. R. Williamson, T. F. Quatieri, B. S. Helfer, R. Horwitz, B. Yu, and D. D. Mehta, “Vocal biomarkers of depression based on motor incoordination,” in *Proceedings of the 3rd ACM International workshop on Audio/visual emotion challenge*. ACM, 2013, pp. 41–48.
- [52] A. J. Larrazabal, N. Nieto, V. Peterson, D. H. Milone, and E. Ferrante, “Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis,” *Proceedings of the National Academy of Sciences*, vol. 117, no. 23, pp. 12 592–12 594, 2020.