Multimodal Depression Severity Prediction from medical bio-markers using Machine Learning Tools and Technologies

Shivani Shimpi[1]**, Shyam Thombre**[2]**, Snehal Reddy**[3]**, Ritik Sharma**[1]**, and Srijan Singh**[1]

[1]University of Mumbai
[2]Indian Institute of Technology, Bombay
[3]Indian Institute of Technology, Kharagpur

ABSTRACT

Depression has been a leading cause of mental-health illnesses across the world. While the loss of lives due to unmanaged depression is a subject of attention, so is the lack of diagnostic tests and subjectivity involved. Using behavioral cues to automate depression diagnosis and stage prediction in the recent years has relatively increased though the absence of labeled behavioral datasets and a huge amount of possible variations prove to be a major challenge in accomplishing the task. This paper proposes a novel Custom CM Ensemble approach and focuses on a paradigm of cross-platform smartphone application that takes multimodal inputs from a user through a series of pre-defined questions, sends it to the Cloud ML architecture and conveys back a depression quotient, representative of its severity. Our app estimates the severity of depression based on a multi-class classification model by utilizing the language, audio, and visual modalities. The given approach attempts to detect, emphasize, and classify the features of a depressed person based on the low-level descriptors for verbal and visual features, and context of the language features when prompted with a question. The model achieved an accuracy of 91.56% and a precision value of 88.46%. Further optimization reveals the intramodality and intermodality relevance, through the selection of the most influential features within each modality for decision making.

1 Introduction

According to the statistics from the World Health Organisation (WHO), around 800,000 people commit suicides every year due to improper anxiety management or overlooked depression. It is to be noted that only one out of five people receive the treatment that is consistent with the current practice guidelines. When analyzed it was found that in about 74% of mental health cases, patients have had their first outbreak before the age of 24. Recent research from WHO stated that on an average there are about 20 suicide attempts for every suicide committed, providing enough room for rescue. Additionally, the epidemiological data suggests that the prevalence and severity of mental health problems are escalating in volumes. While with respect to the current situation, the Lancet review produced a psychological report suggesting the impact of quarantine to be wide-ranging, substantial, and long-lasting; the 24 out of 3166 studies clearly linked the longer duration of quarantine to poorer mental health specifically the post-traumatic stress symptoms, avoidance behaviors, anger among people and a lack of empathy in health-workers; most of which suggested a better approach towards mental-health during the quarantine to be just as essential as physical.

The Internet and recent mobile health interventions have proven helpful time and again to overcome problems related to social stigma and accessibility. These services have been implemented to supplement existing mental health treatments or to expand the limited access to quality mental health services while simultaneously being matched with 70% of patients showing interest in using mobile apps to self-monitor and self-manage their mental

health. Applications like Woebot, Wysa, Replika, and Ginger.io have proven that digital interventions for anxiety and depression detection bio-markers have an empirical support with outcomes comparable to therapist-delivered psychotherapies or cognitive-behavioral therapies (CBT).

There are two methods to digitize the mental screening of a user - The first approach models the outcome based on language responses of the user relevant to the questions, while the second approach focuses more on non-verbal features which would be the low-level signal descriptors including prosodic, glottal, and spectral features from the audio modality, and the muscle and micro-muscle movements are encoded in the form of action units (AUs) based on the Facial Action Coding System (FACS) from the visual modality.

2 Literature Survey

While Gong et al.[1] developed an ensemble of language, audio, and visual modalities based on the type of questions, Yang et al.[2] utilized deep learning frameworks to combine these multiple modalities conditioned on manual question selection. A key reason to fuse all the three modalities is supported by Morales and Levitan[3] where they tested speech versus text features, further concluding that a multimodal system leads to a better performance over its unimodal equivalent.

2.1 Language

Most of the depression detection studies in literature are usually social media context analysis[4] at most combined with visual representations[5]. In rare instances, conversational text generated from clinical interviews is examined on its own, commonly investigated in terms of semantic or linguistic features coupled with audio or visual modalities. Since a clinical conversation involves questions and answers from different parties, semantic information is separated into content and context analysis. Word representations using Global Vectors for Word Representation (GloVe), following a high-level feature learning method concludes that the semantic analysis of dialogue scripts via text-based features is the most promising depression detection method with respect to other modalities[6]. Pampouchidou et al.[7] extracted features from the verbal communication recorded in the transcript file provided in the data. A depiction in comparison[8] for different ways of modeling a conversation with the combination of audio and text features in specific, they compared context-dependent, context-free, and sequence modeling methods using Word2Vec as their word-embeddings. The authors adopt clustering techniques[9] to quantify anxiety and depression indices on questionnaire textual data. Further, the correlation amongst anxiety, depression, and social data is investigated. For most of social media data, the text analysis is performed on short text, although, these classifiers do not show similar results in a conversational setting which happens during the counselling/screening sessions.

2.2 Audio

Inspired by the present deep learning performance in detecting emotion, a deep learning based approach using Convolutional Neural Networks (CNN) on Audio and Deep residual networks (ResNet) was implemented by Tzirakis et al.[10]. They claimed that their approach gave the best result of the time. Speech being non-invasive and non-intrusive is suitable for nearly all automation tasks. The characteristics like diminished prosody, the amount of monotonous verbal activity production, and the energy in speech are important bio-markers for distress diagnosis. Cummins et.al.[11] provided an exhaustive review of depression and suicide risk assessment through the usage of vocal bio-markers to associate clinical scores for depressed classes. A distress assessment in speech signals was performed[12] to infer emotional information expressed while speaking which quantified the anger, valence, arousal, and dominance values. Another research[13] provided a comparative study that assisted in depression detection by using the mel-frequency cepstral coefficients (MFCCs) and damped oscillator cepstral coefficients (DOCCs). Cummins et.al.[14] investigated a change in spectral and energy densities of speech signals for depression prediction by analysing the acoustic variability in terms of weighted variance, trajectory of speech signals, and the volume to measure depression.

2.3 Visual

Convolutional multiple kernel learning approaches[15] have been performed for emotion recognition and sentiment analysis tasks in videos. The Facial Expression Recognition and Analysis challenge (FERA 2017) provided the dataset[16] to estimate the head pose movements and identify the action units against them which were required to quantify the facial expressions. A thorough meta-analysis[17] of facial gestures to identify schizophrenia-based event triggers were performed and a coherent link was established[18] between attention deficit hyperactivity disorder (ADHD) and emotional reactions due to the lack of attention. Dalili et.al.[19] conducted a thorough study for meta-analysis of the association between existing facial emotion recognition and depression techniques, whereas, the authors of another research[20] relied on temporal Long Short Term Memory (LSTM) network based techniques to capture the contextual information from videos in a sentiment analysis task. An open-source interactive tool, OpenFace[21], is used to estimate the facial behaviour based on a facial action coding system (FACS). The dataset that we used includes the low-level visual features like face landmark regions, head-pose and eye-gaze estimation values, all of which are converted into a reliable facial action unit extracted using OpenFace.

3 Feature Extraction and Selection

This research is based on the complete Distress Analysis Interview Corpus - Wizard of Oz (DAIC-WOZ) dataset and three samples of depressed class 4 and 5 from the Extended DAIC dataset provided by the University of Southern California. The complete dataset comprises of 192 sessions out of which 189 are the DAIC-WOZ interactions between the patient and an animated virtual interviewer (Ellie). The clinician, through the digital avatar, asks a series of questions specifically aimed at identifying depressive symptoms. The dataset has about 50 hours of interviews in total each varying from 7-33min (with an average of 16min), each session including language, audio, and visual modalities in the form of transcript of the interaction, participant audio files, and facial features respectively. Since the language modality in the dataset was manually labeled, it ensures a good quality of the transcripts. The interactions are based on an internationally accepted 8 question-based **P**atient **H**ealth **Q**uestionnaire. The labels provided are the PHQ-8 scores ranging from 0-24 alongside the binary class values for every data point. Table 1 includes the data splits. We have further bifurcated these scores into 5 classes where Class 0 ranges from 0-4, Class 1 is from 5-9, Class 2 is 10-14, and so on. A few major issues we faced with the dataset are as follows:

- Insufficient data - The provided data consists of very less data points and does not suffice for improvising on the evaluation metrics for every modality causing model overfit.

- Huge Class Imbalance in the dataset - The data being very less, it was highly class imbalanced. Class 0 had 64 samples, while class 4 consisted only of 5 samples making it even more challenging to show depressed classes to the model.

- Weakly Annotated Data - Even though we have enough data for a single participant, just one label is associated with a huge set of features, or in other words, although the depression score and PHQ8 scores are correlated, one characteristic does not necessarily predict another.

3.1 Language

The dataset provides a detailed transcript of the whole interview obtained from the audio using a Transcriber and was manually edited to improve the quality of translation. The entire interview has a specific label, so it is not segmented based on time slots. The sequential modeling approach we took is a context-free method that processes patients' response without any regard to dialogue-based interactions. The pre-processing of text followed certain rules that needed every letter to be lowercase, unaccented conversions of the English words, and ASCII conversions for the foreign words. Disfluencies in speech, for example, hesitation were kept raw in the transcript, while the verbal metadata, for example, laughter, sigh, coughing are put in angled brackets *<laughter>*, *<sigh>*, *<cough>*. In sequence modeling, we removed the questions from the interviewer and considered all the answers from the patient extended into a list as one single data-point. Every sample was lemmatized to get a specific word density associated

with a depression severity class value. Lemmatization was preferred over stemming to retain the knowledge of the context and the discrimination between words having different meanings depending on the part of speech.

3.2 Audio

The dataset consisted of raw audio (.wav format) files of the interview for all the participants which contained the voice of the participant as well as the interviewer for the duration of an entire interview. It also included the features extracted from these audio files using the open-source **Co**llaborative **V**oice **A**nalysis **Rep**ository (COVAREP) software. These features had the frame-level of 20 milliseconds window and a 10 milliseconds shift for an entire recording of the interview comprising of-

- spectral features - Mel-Frequency Cepstral Coefficients (MFCCs) 0-24, Harmonic Model and Phase Distortion Mean (HMPDM) 0-24, Harmonic Model and Phase Distortion Deviation (HMPDD) 0-12.

- prosodic features - voicing probability, pitch, and the first five formants.

- voice quality - maximum dispersion quotient, normalized amplitude quotient, difference in amplitude of the first two harmonics of the differentiated glottal source spectrum, quasi open quotient, parabolic spectral parameter, spectral tilt/slope of wavelet responses, and shape parameter of Liljencrants-Fant model of the glottal pulse dynamics.

All these features were normalized with zero mean and unit variance.

3.3 Visual

Although the relationship between the linguistic context-based models and mental illness level are more prominent, the visual features embody a key role to reinstate the correlation of facial expressions to depression. The visual modality of the dataset has a lot of features, including a complete 3D and 2D facial points, action coding units, pose and gaze, etc. that are extracted using OpenFace. Since we did not have sufficient data points and a large number of features for each data point, the critical factor was to ensure a thorough feature engineering for the models. The patients who suffer from depression often have distorted facial expressions, for example, eyebrow twitching, dull smile, frowning faces, aggressive looks, the tendency to lower their head more as compared to the non-depressed people, restricted lip movements, and reduced eye blinks. OpenFace is a widely popular interactive tool to estimate the facial behaviour. In this paper, we have used the low-level visual features extracted using OpenFace that included-

- 68 2-dimensional and 3-dimensional points on the face.

- regression outputs for each action unit and binary labels reflecting the presence of a specific action unit.

- gaze as an output of 4 vectors, first two vectors in world coordinate space describing the gaze direction of both eyes, while the other two vectors describing the gaze in head coordinate space, for example, if the eyes are rolled up, the vectors will indicate up even if the head is turned or tilted.

- Felzenswalb's Histogram of Gradient (HoG) binary format files aligned on the 112x112 area of the face, resulting into a 4464 vector per frame stored in a byte stream.

- Pose as an output of position coordinates (X, Y, Z) and head rotation coordinates (R_x, R_y, R_z). Position is in the world coordinate system in milimeters while rotation is in radians and in Euler angle conventions to get a proper rotation matrix the product of the head rotation coordinates is taken.

4 Methods

Automatic depression detection research can either predict the classification results or a severity score, to associate with the mental state label and PHQ-8 score. The approach we took considers using the language, audio, and visual modalities working interdependently with each other.

Split	Depressed	Non-Depressed	Σ
Training	30	77	107
Testing	14	33	47
Development	12	23	35
Σ	56	133	189

Table 1. DAIC-WOZ Dataset Overview

4.1 Model Development

4.1.1 Language

Presence of depression is a binary and multi-class classification model, predicting healthy/depressed state of the user. Categorical accuracy metric is presented as the performance score. Due to the lack of data, the model complexity was kept relatively simpler. We created dense word encoding using an embedding layer and averaged the outputs of each feature map through Global Average Pooling. In order to avoid overfitting, the total number of parameters in the model were reduced. This was then connected to a Dense layer with ridge (L2) regularizer. Another variation of the model was the introduction of Bidirectional Long Short Term Memory (Bi-LSTM) Networks with dropout regularization used to access both the preceding and succeeding context representations connected to a Dense layer with ridge (l2) regularization, to avoid the problem of vanishing gradients we used the leaky Rectified Linear Unit (ReLU). Since neither of the models provided a satisfactory result based on the amount of data we had, using a pre-trained model was better option. We implemented DistilBERT[22] which is a lighter version of Bidirectional Encoder Representations from Transformers (BERT). DistilBERT compared to the original BERT base has up to 40% less parameters than the uncased BERT and runs 60% faster while preserving over 95% of BERT's performances as measured on the GLUE language understanding benchmark.
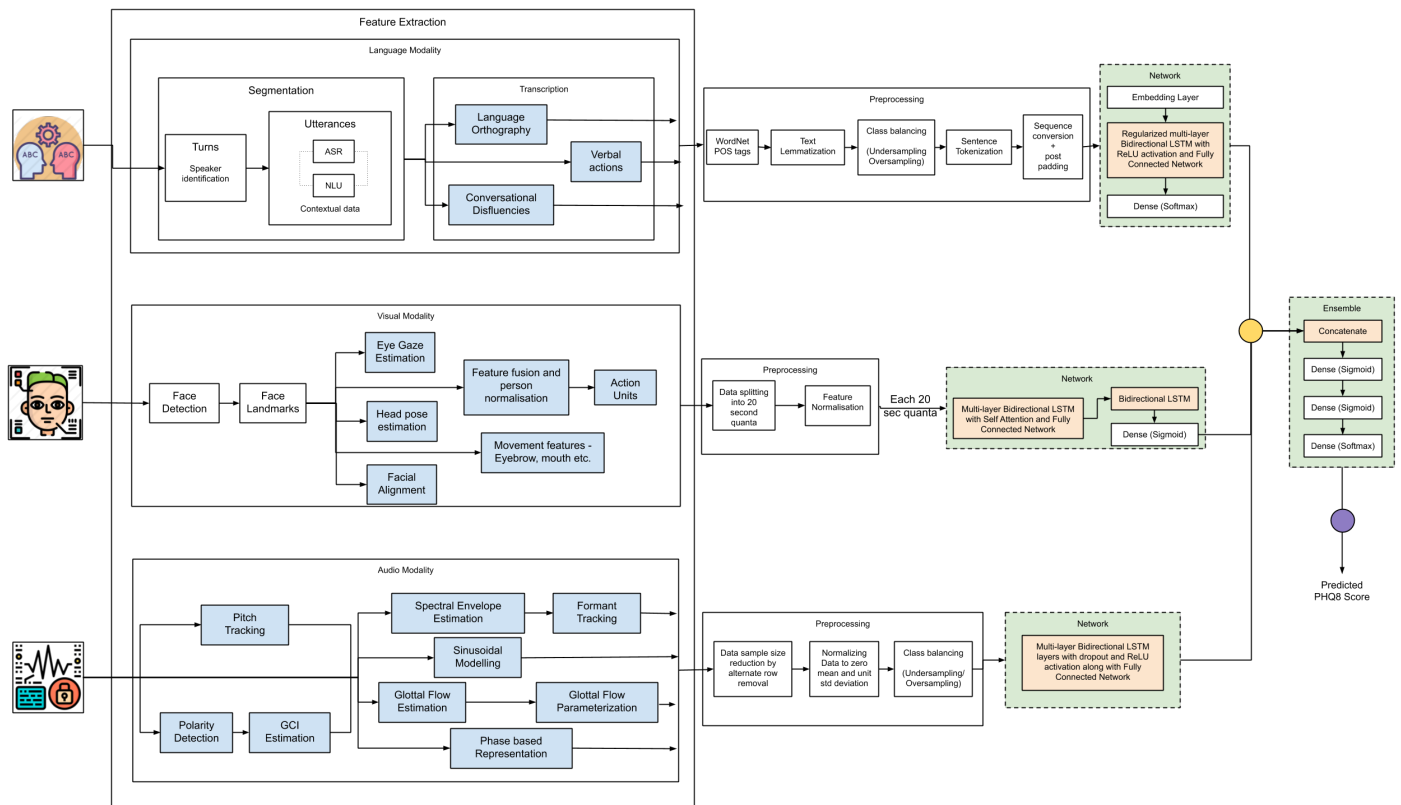

Figure 1. Multimodal Pipeline

4.1.2 Audio

The approach to audio model can be divided into binary classification and regression. Binary classification has the split label of 10 for positive and negative classes and the regression performed on the given scale of 0-24. All the models were trained using the Stochastic Gradient Descent (SGD) optimizer and Adam Optimizer. For the regression models, SGD optimizer seemed to perform better, whereas for binary classification, the performance for both of them was very similar. We also tried weight decay along with dropouts for all the models. Although this helped to improve the performance in some of the cases, it was not very useful in the most.

In order to resolve the issue of class imabalnce, we balanced all the classes using the undersampling and oversampling techniques. The binary classification models developed using GRU and fully connected layers showed high signs of overfitting with a random probability that scaling down the weights with regularization did not seem to resolve, so another way to approach the problem was by changing the weights distribution for which the dropout units were used. The GRU outputs were then extracted post training, and were provided to train a Support Vector Machine (SVM) using Gridsearch for the hyper-parameter tuning. In another approach, we replaced the Support Vector Machine with Random Forest Classifier. A significant overfit in models was consistent throughout the approaches we tried, thus we decreased the model complexity by reducing the number of input features through feature selection using MFCC, HMPDM, and HMPDD to ensure a spatial relationship as features for the regression model using Temporal Convolutional Networks (TCN) connected to dense layers with and without dropout regularization. We then created a Bi-directional LSTM network connected to dense layers with dropout regularization. In order to resolve the lower dataset issue, we segmented the audio clip at various lengths and replicated the label for each segment which increased data set and also reduced input sample size.

	Feature Learning	Supervised Training	Testing
Classic Deep Learning	Audio	Audio	Audio
Multimodal Fusion	L+A+V	L+A+V	L+A+V
Cross Modality Learning	L+A+V	V	V
Custom CM Ensemble *(Our approach)*	L+A+V	L+A+V	V

Table 2. Multimodal Learning settings where L+A+V refers to the Language, Audio, and Visual modalities

4.1.3 Visual

Logisitic classification was performed on the entire dataset using all the features. Since it cannot handle time series data format, the average chunk of the data had to be fed to the model and the average output for the entire video was the final output which led to a model overfit that did not get resolved by a reduction in model complexity and regularisation. The 3D data points were considered not too helpful for the model when the features were sorted based on importance using XGBoost. We created more features with reference[23] from the 2D facial landmarks:

- Distance between eyelids (vertical distance) and horizontal width of the eyes.

- Distance between lips (vertical distance) and horizontal width of the mouth.

- Motion of eyebrows, eyes, mouth and head using the previous frames.

Redundant features like 3D facial landmarks were eliminated and the other features were processed using feature engineering. The visual model is bi-stacked vanilla LSTM with Global Max Pooling layer in the end followed by a feed-forward neural network. The entire video was fed as a single data point which led to an increase in the prediction accuracy, but lack of data abundance was an imminent problem for this approach, since the number of interviews were of the order of 150. In order to resolve the issue of the amount of data points, each video was partitioned into 20 second intervals, which were then treated as a single data point. This led to a considerable increase in the performance resulting in no signs of model overfit opening the room to increasing the model complexity.

The approach consisted of using the Bi-LSTM network and increasing the number of dense layers. Since the feature set we used were diverse with pose, gaze, facial action unit coded values, 2D facial landmarks etc., a good

way for the model to learn from the features was through an attention layer that allowed the model to learn to "attend to" a certain set of the vector which was more useful in the ongoing context to enable easier and higher quality learning.

4.1.4 Ensemble

The intention for a multimodal approach and combining the three modalities was to produce better predictions as compared to the individual models. For this purpose we made use of the learning settings[24] mentioned in Table 2. The classical deep learning techniques were used to develop individual models for each of the three modalities. The first approach towards the multimodal learning was the multimodal fusion which is the typical setting in the prior work, in this paper, we propose the custom CM learning setting for depression severity prediction based on the multi-class classification objective.

1. **Multimodal Fusion:** We trained the visual model separately, further removing the final classification layer and used this an input for the final ensemble model. The visual pipeline takes in the 20 second interval clips and then performs feature extraction. We clip the final few layers of the pipeline and then pass the features to the model thus forming the latest representations of the 20 second clips. These are then concatenated and passed to a Bi-LSTM network to form the visual representation of the entire video. The audio and text pipelines generate a one dimensional latent space representation for the entire interview. The final language, audio, and visual representations are then concatenated, which could then be passed to dense layers to form the basis of our prediction. This approach has an important underlying assumption that the individual models have learned good representations of the interview for the corresponding modalities, which has been justified by the 80% accuracy for individual language and audio modalities, and 88.24% accuracy for the visual modality. The depiction for the multimodal pipeline can be seen in the Figure 1.

2. **Custom CM Ensemble:** The performance of the visual model was much better than the other individual models, so we wanted to evaluate if providing all three modalities would help improvise the model and help it learn better representations. We got the representation from the visual model by providing the language, audio, and visual modalities as present during feature learning phase.

4.2 Model Deployment

The data was recorded from an Android or iOS device using our application. This was given as an input to the cloud where the machine learning models predicted the results and sent them back to the application. The models were deployed on the Google Cloud platform using serving input functions where the model could make predictions based on the REST API call requests from the client side. An abstract workflow of the application is depicted in Figure 2 The two main components of the application development are the front-end and back-end discussed in detail below.

4.2.1 Front-end

React-Native with Redux as the state management tool. Node Package Manager (NPM) was used to manage the JavaScript packages used for the UI components.

1. **React-Native:** It is an open-source framework based on the JavaScript library which was used to design the front-end of the cross-platform Android and iOS application.

2. **React-Redux:** Since the app involves a lot of components, state changes, back-end data fetch, caché storage, etc., the state management gets challenging. At some point, it is difficult to understand what happens in the application as the control over the state functionality is lost, also when the system is opaque and non-deterministic, it is hard to reproduce bugs or add new features. Thus, we need an application state management technique for which we used React-Redux making it easy to store caché and reuse state components instead of loading them multiple times.
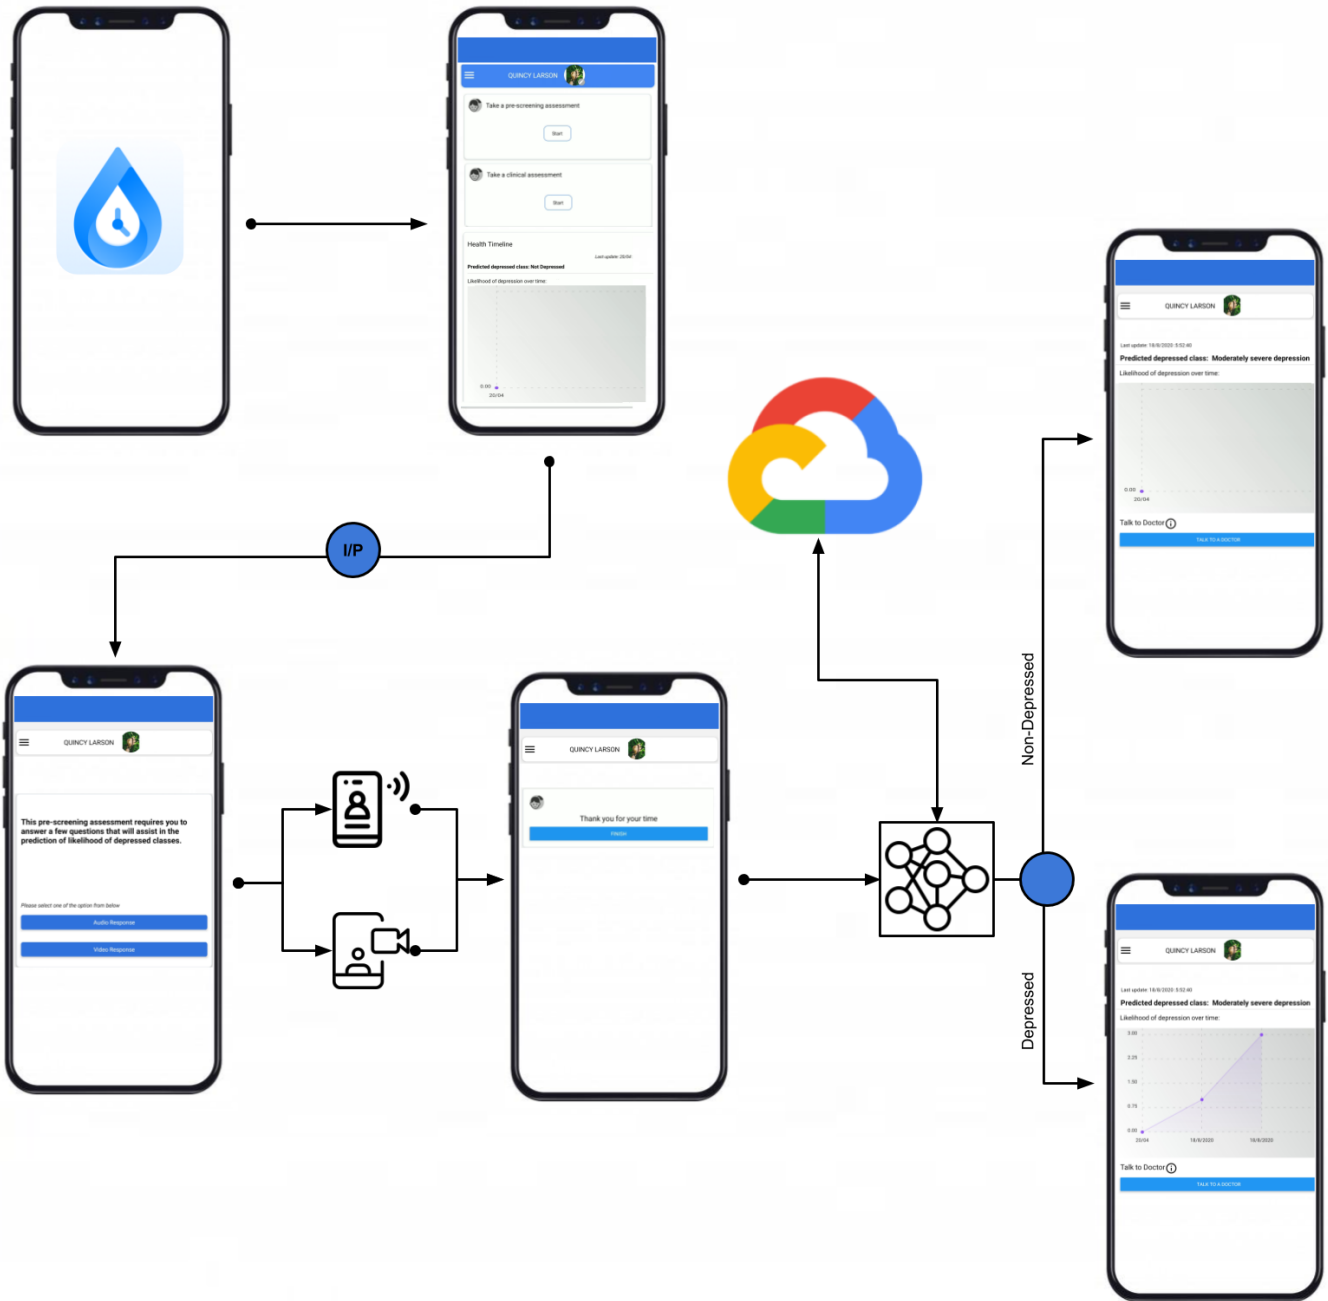
Figure 2. User Workflow for the application

3. **Node Package Manager:** We use the Node Package Manager (NPM) to manage the JavaScript packages required for User Interface (UI) components of the application. Out of over 20 packages that we used, the main six packages are listed below:

    (a) The @react-native-community/google-signin is used for signing into the application using the Google account.

    (b) Firebase services:

        i. Firebase-admin was used for user management.

    ii. The @react-native-firebase/app is used to register the application to the Google Cloud Platform.

    iii. The Firebase authentication @react-native-firebase/auth is used for user login creation that allows the login by using a custom email address or google credentials.

    iv. Firestore @react-native-firebase/firestore database is a server-less live synchronised powerful query engine with offline mode that is used for storing the user data.

    v. The cloud messaging service using @react-native-firebase/messaging enables a messaging service for the user and doctor.

    vi. The firebase storage @react-native-firebase/storage is mainly used for storing the user responses that permits file storage. It not only stores user response. It also stores some applications files, user profile data like profile images and background images along with the ML models.

(c) React-navigation: This is used for navigating and passing the data between multiple pages.

    i. @react-navigation/native: React Navigation is made up of some core utilities which are used by navigators to establish the navigation structure in the application.

    ii. @react-navigation/drawer: This is used for drawer navigation in the application user interface.

    iii. @react-navigation/stack: The stack navigation provides to transition between screens where each new screen is placed on top of a stack.

(d) React-Native-video-recorder is used to record videos for the user answers.

(e) React-Native-audio-recorder is used to recording audios for the user answers.

(f) React-Native-webRTC is used for a video calling feature in the application.

4.2.2 Back-end

Most of the back-end was structured using the Google Cloud Platform (GCP) involving services like authentication, storage, database, and containerizing the models. A detailed overview of the application deployment and workflow is depicted in the Figure 3.

1. **Authentication:** It involves back-end management of the user login and the sign-up services, while also ensuring the creation, modification, and deletion of a user. Authentication was done via Firebase-Authentication from the Google Cloud Platform (GCP) using the @react-native-firebase/auth from React-Native.

2. **Storage:** It involves storing the user files like the patient reports, medication, images, video, audio, history, and logs files. This service was implemented with the help of Cloud-Storage (Firebase) from the Google Cloud Platform (GCP). The @react-native-firebase/storage dependency from React-Native helped to store files from the application to back-end storage.

3. **Database:** We used Cloud Firestore (Firebase), a scalable cloud database to store all the data required for both the front-end and back-end, using its client library for both Python v3.7 and Node.js to interact with the database. It involves the database for the user profiles for the patients and doctors another database that purely comprised of the results. The Cloud-firestore (firebase) from the Google Cloud Platform (GCP) was used for database and @react-native-firebase/firestore dependency from React-Native helped data flow between the application and back-end.
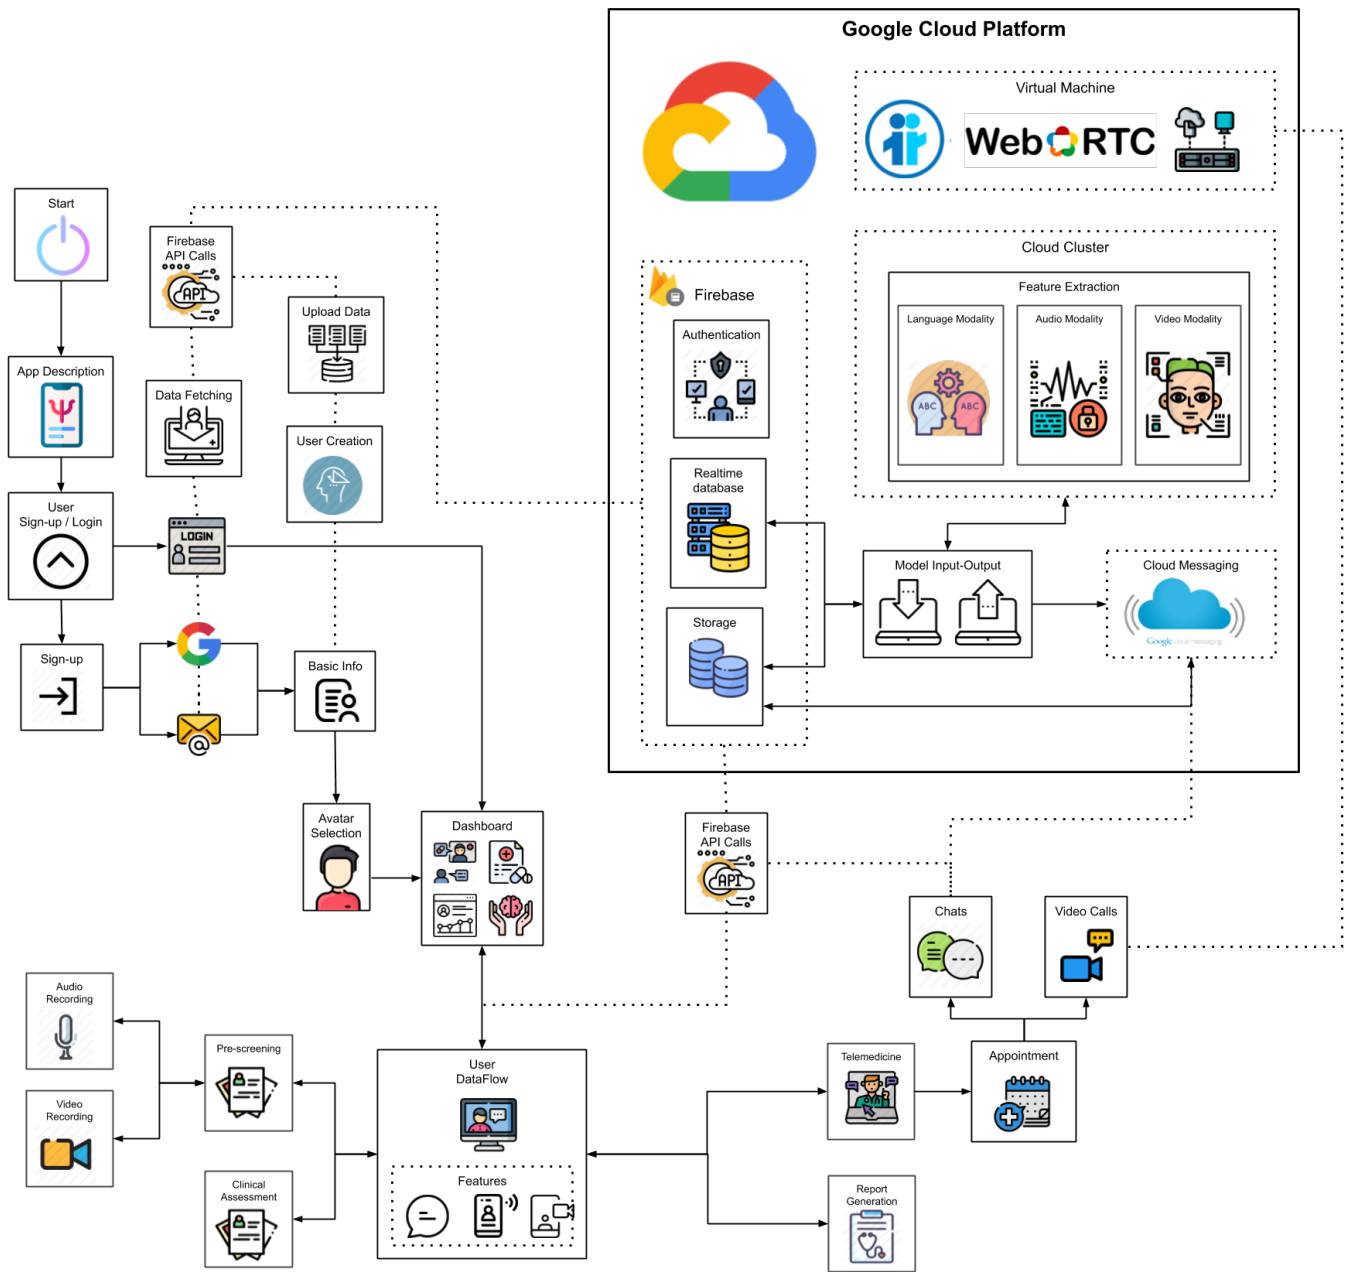
Figure 3. Detailed Application Workflow

4. **Cloud Function:** Cloud Functions (Firebase) with Python v3.7 as runtime was used to implement the file transfer feature in the chat messaging part of the application using the @react-native-firebase/messaging dependency from React-Native.

5. **Models:** It was developed using Python programming language v3.7 (open-source). The models were containerized using Docker (open-source) into separate microservices for better isolation, scalability, and deployment. These microservices were then deployed on Google-Cloud-Run (GCP) with GoogleKubernetes-Engine (GCP) for container orchestration. Gunicorn WSGI server (open-source) was used to expose these microservices as RESTful APIs and Google-Cloud-Tasks (GCP) for asynchronous inter-service communication.

5 Results and Conclusion

The class imbalance in dataset was handled by the undersampling and oversampling techniques; for undersampling, all the samples corresponding to depressed labels were used and the same number of samples of non-depressed labels were randomly chosen from the remaining set, whereas, for oversampling, all the samples corresponding to the depressed and non-depressed labels were used and the samples of depressed labels were reused to match the number of non-depressed labels. All the models and results are mentioned in the Table 3.

5.1 Language

While the training loss as shown in the Figure 4 starts to decrease signifying that the model is learning well on the training data, the increase in test loss signifies that it failed to make any progress on the test data. A key thing to note here is the accuracy curve for the train and the test set - the training accuracy shows an increase along with the test accuracy, meaning that the model predicted some classes with confidence but not the others. Loss generally measures a difference between raw prediction that could be a float value and classes, while accuracy measures difference between threshold prediction which signify the classes and available classes. If the raw predictions change, the loss changes but the accuracy is more resilient because the predictions need to go over or under a specific threshold to be able to change the accuracy. The curves that we obtained by training the language model signified -

- Noisy Labels: This was possible from comparing the false predictions when the test loss is minimum and test accuracy is maximum. A solution to that would have been to check if the samples were correctly labelled but that is temporarily out of the research scope.

- Feature relevancy: It is probable that the model did not have enough aspect of information to be certain about a specific class. A better way to resolve this could have been through experimenting with more and larger hidden layers, which when we tried resulted into model overfit, so the alternative approach is to use a much larger dataset with more severely depressed class values provided with same features.

- Biased Class Learning: This means that the network at a given epoch might severely overfit on some classes but still be learning on the other classes, this is the cause of the highly imbalanced classes in the dataset.
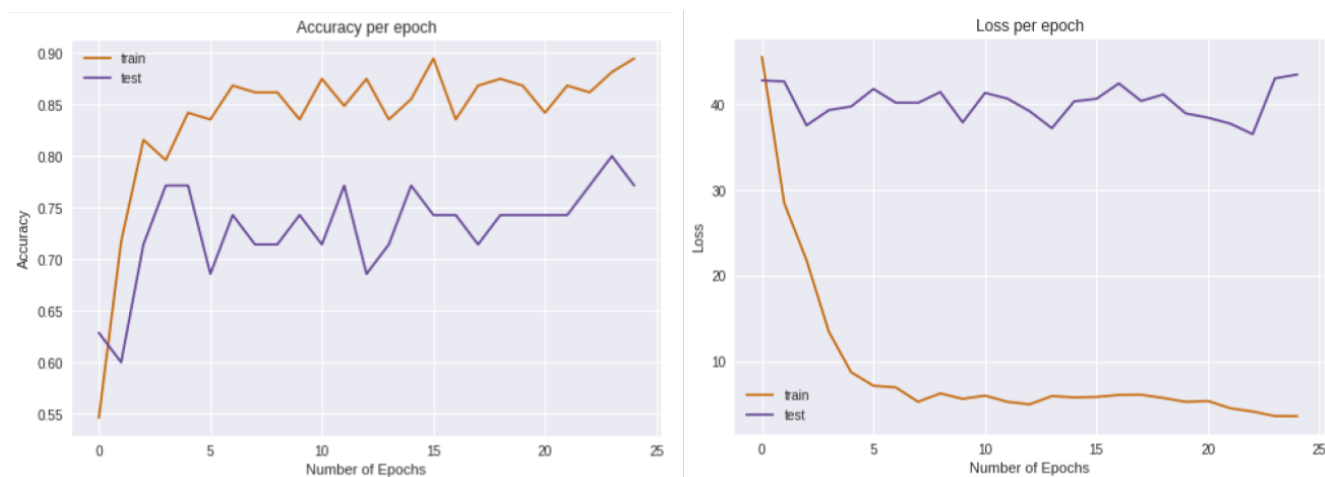

Figure 4. Performance curves for Language model

Since the loss is stable and the accuracy increases it implies that the model is making good use of the regularization techniques that we used. However, it stops learning or making any significant progress on the test dataset and is stable as depicted in the Figure 4 for performance curves for the language model.
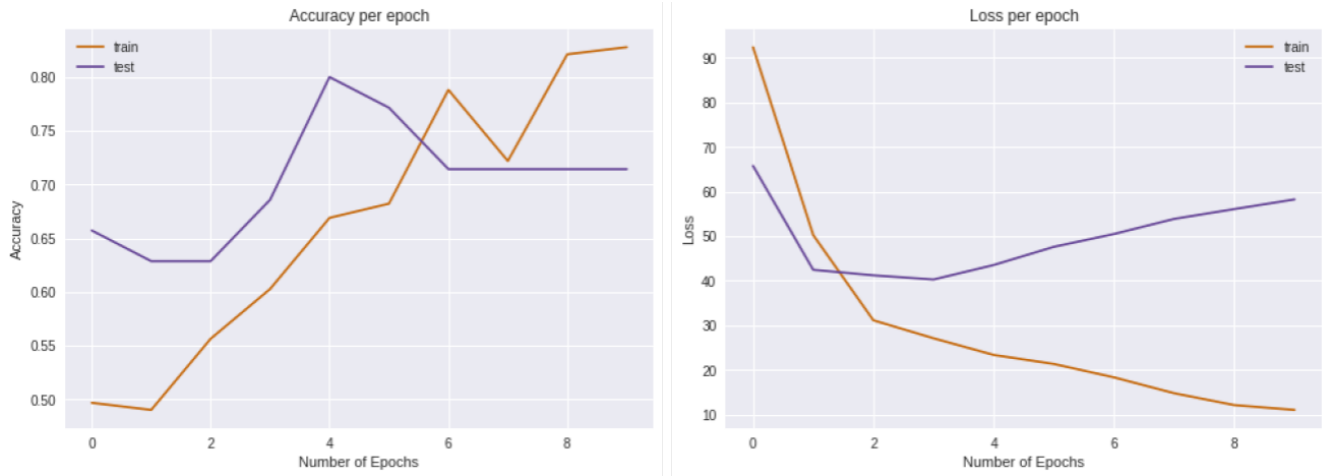
Figure 5. Performance curves for Audio model

5.2 Audio

The balanced class techniques significantly improved the LSTM model for binary classification and regression, achieving an accuracy of 80% and an F1-score of 0.63 with a 6.1 RMSE error, although these models did not address the multi-class classification problem.

Model	Modality	Processing	Accuracy	Precision	Recall	F1-Val
TF-IDF with Logistic Regression	Language	Preprocessed	65.71%	0.57	0.44	0.533
TF-IDF with Random Forest	Language	Preprocessed	<50%	-	-	0.42
CountVectorizer with Logistic Regression	Language	Preprocessed	63%.	-	-	0.43
Embedding + GlobalAveragePooling + FC	Language	Lemmatized	71.96%.	-	-	-
Embedding + Bi-LSTM + FC + ReLU	Language	Lemmatized	65.71%.	-	-	-
Embedding + dropout Bi-LSTM + ReLU	Language	Lemmatized	62.86%.	-	-	-
Embedding + reg Bi-LSTM + FC + ReLU	Language	Lemmatized	65.71%.	-	-	-
DistilBERT	Language	Lemmatized	71.96%.	-	-	-
GRU + FC	Audio	COVAREP	51	0.31	0.33	-
GRU + SVM (With/Without SMOTE)	Audio	COVAREP	64.71%	0.5	0.58	-
GRU + Random Forest + FC	Audio	COVAREP	70.58%	0.63	0.42	-
GRU + Random Forest (SMOTE) + FC	Audio	COVAREP	67.65%	0.55	0.5	-
SVM ± Feature Selection (SMOTE)	Audio	Preprocessed	63.63%	0.5	0.58	-
RF + Feature Selection (SMOTE)	Audio	Preprocessed	63.63%	0.5	0.5	-
Random Forest (SMOTE)	Audio	Preprocessed	60.6%	0.42	0.25	-
TCN + Random Forest (SMOTE)	Audio	COVAREP	77.14%	0.7	0.58	-
Logistic Classification	Visual	OpenFace	57.9%	-	-	-
Vanilla LSTM (entire video)	Visual	OpenFace	64%	-	-	-
Bi-LSTM (chunks)	Visual	OpenFace	79%	-	-	-
Bi-LSTM + Self Attention (chunks)	Visual	OpenFace	88.24%	-	-	-
Multimodal Fusion	AVL*	Preprocessed	73.53%	-	-	-
Custom CM Ensemble	AVL*	Preprocessed	91.56%	0.8846	0.8518	0.8678

Table 3. Models and corresponding evaluation metrics, where *AVL signifies the audio, visual, and language modalities respectively

The Synthetic Minority Oversampling Technique (SMOTE) was tried on the latent space representations that were learned by Gated Recurrent Unit (GRU), Long Short-Term Memory (LSTM), Temporal Convolutional Network (TCN) classified using Support Vector Machines (SVM) and Random Forest. However this improved the TCN model with about an 80% binary accuracy, it did not perform very well on the others. Due to the smaller test set, the test accuracy curve was observed to fluctuate a lot. The dataset size seemed to be the main issue that needed to be resolved for which a pre-trained for sentiment or emotion analysis (transfer learning) could have provided a better solution. As a workaround to the weak annotations, splitting the audio of an interview and replicating the labels generated some noisy labels, worsening the issue of class imbalance and it eventually did not help in improving the overall model performance. An increase in the loss curve as depicted in Figure 5 can be resolved with better feature representations and a larger audio dataset.
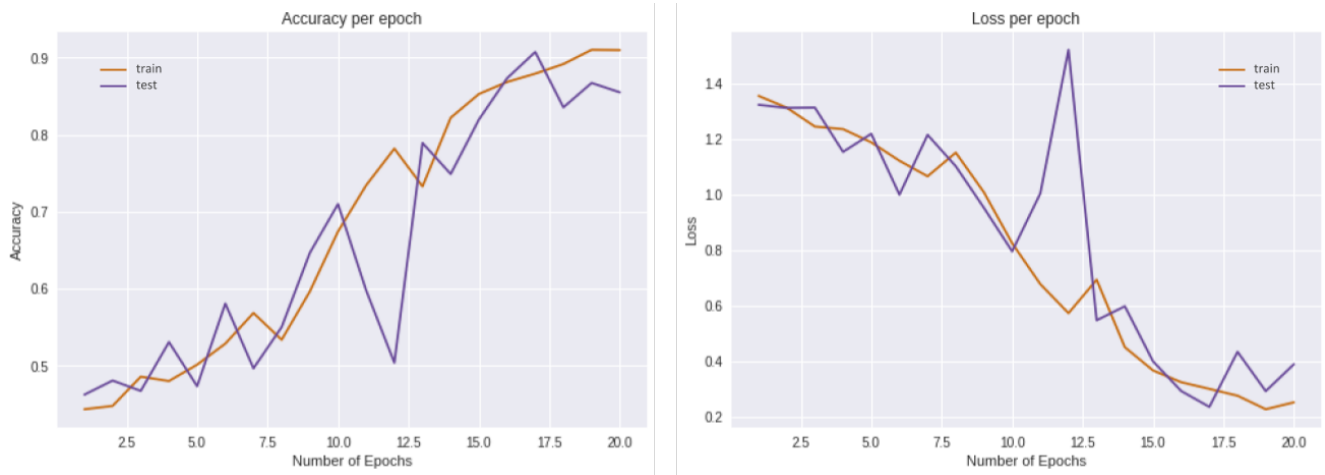

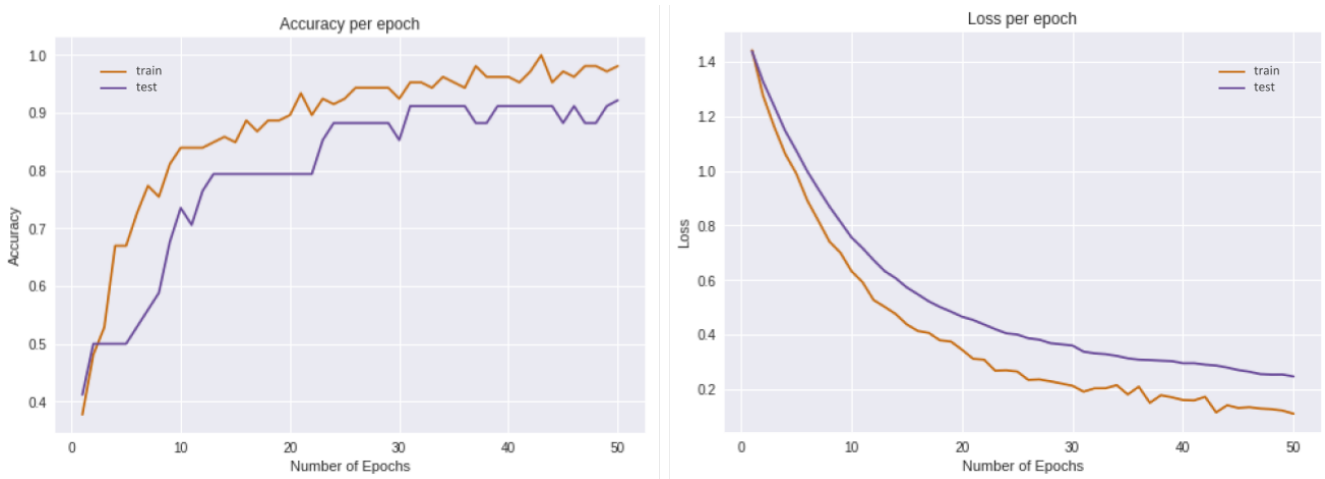**Figure 6.** Performance curves for Visual model

Figure 7. Performance curves for the custom CM ensemble model

5.3 Visual
The visual model performed significantly well in terms of predicting depression severity in users. The use of attention mechanism in the model led to a positive increase in the evaluation metrics for a batch size of 16, 20 epochs, and 376 steps per epoch. The decreasing train and test loss curves as depicted in Figure 6 imply that the model learned from the train as well as the test set. The sudden drop for one of the epochs was due to the imbalanced multi-class dataset, where the model could not very confidently predict the severely depressed classes.

5.4 Custom CM Ensemble

The multimodally fused model based on concatenation of the pre-final layer of the LSTM outputs degraded the performance when compared to the individual models. The graph for the multimodal fusion is depicted in the Figure 6. The Custom CM Ensemble model learned much better representations as compared to multimodal fusion. The graph for the Custom CM Ensemble model as depicted in Figure 7 shows a smoother loss curve signifying a better representation learning on all the modalities, making it a better fit. In conclusion, improvising the visual model using the Custom CM Ensemble approach provided better predictions. Based on the loss and accuracy curves, it is safe to conclude that if provided with sufficient dataset, the model will predict more accurate and precise multi-class predictions than the current ones.

6 Discussions

Potential approaches include changing the training algorithm to One-Shot learning, and using the Siamese Neural Network to address the problem of a smaller dataset and class-imbalance, thus following an approach of features comparison rather than feature representation. There is also scope for some changes in the approach for feature extraction in the audio model using the COVAREP toolbox, by pre-processing the audio files so as to try and emphasize on the more important speech parameters using signal processing and denoising. A significant decrease in the test loss per epoch with an improvised test accuracy for the visual model was observed making it safe to conclude that the model continued to learn from the provided input features. The language model can be improvised by using the pre-trained text embeddings and the original base model for BERT and a dialogue-based approach with One-Shot learning. Multiple Instance Learning in terms of multi-class classification can prove to be helpful to deal with the problem of smaller dataset.

7 Author Contributions

Sh.S. conceived the study, devised the project plan, worked on the development of models for the language modality of the experiments, programmed the Custom CM Ensemble model, and wrote the manuscript with input from the other authors. Sh.T. contributed to study conception, developed the models for audio and language modalities, and contributed to the manuscript. Sn.R. programmed the models for visual modality and contributed to the manuscript. R.S. conceived the application development, developed the front-end and back-end of the prototype/application, and wrote the manuscript. Sr.S. worked on the front-end of the application, deployed the ML models in the cloud, and setup the back-end.

8 Acknowledgements

References

1. Gong, Y. & Poellabauer, C. Topic modeling based multi-modal depression detection, DOI: http://doi.acm.org/10.1145/3133944.3133945 (2017).

2. Yang, L. *et al.* Multimodal measurement of depression using deep learning models, DOI: https://doi.org/10.1145/3133944.3133948 (2017).

3. Morales, M. R. & Levitan, R. Mitigating confounding factors in depression detection using an unsupervised clustering approach (2016).

4. Trotzek, M., Koitka, S. & Friedrich, C. M. Utilizing neural networks and linguistic metadata for early detection of depression indications in text sequences, DOI: https://doi.org/10.1109/TCSS.2019.2894144 (2020).

5. Gui, T. *et al.* Cooperative multimodal approach to depression detection in twitter, DOI: https://doi.org/10.1609/aaai.v33i01.3301110 (2019).

6. Williamson, J. R. *et al.* Detecting depression using vocal, facial and semantic communication cues, DOI: https://doi.org/10.1145/2988257.2988263 (2019).

7. Pampouchidou & Anastasia. Depression assessment by fusing high and low level features from audio, video, and text, DOI: https://doi.org/10.1145/2988257.2988266 (2016).

8. Hanai, T. A., Ghassemi, M. & Glass, J. Detecting depression with audio/text sequence modeling of interviews, DOI: https://doi.org/10.21437/Interspeech.2018-2522 (2018).

9. Hao, F. *et al.* Providing appropriate social support to prevention of depression for highly anxious sufferers, DOI: https://doi.org/10.1109/TCSS.2019.2894144 (2019).

10. Tzirakis & Panagiotis. End-to-end multimodal emotion recognition using deep neural networks, DOI: https://doi.org/10.1109/JSTSP.2017.2764438 (2017).

11. Cummins, N. *et al.* A review of depression and suicide risk assessment using speech analysis, DOI: https://doi.org/10.1016/j.specom.2015.03.004 (2015).

12. Stasak, B., Epps, J., Cummins, N. & Goecke, R. An investigation of emotional speech in depression classification, DOI: https://doi.org/10.21437/Interspeech.2016-867 (2020).

13. Mitra, V., Tsiartas, A. & Shriberg, E. Noise and reverberation effects on depression detection from speech, DOI: https://doi.org/10.1109/ICASSP.2016.7472788 (2016).

14. Cummins, N., Sethu, V., Epps, J., Schnieder, S. & Krajewski, J. Analysis of acoustic space variability in speech affected by depression, DOI: https://doi.org/10.1016/j.specom.2015.09.003 (2015).

15. Poria, S., Chaturvedi, I., Cambria, E. & Hussain, A. Convolutional mkl based multimodal emotion recognition and sentiment analysis, DOI: https://doi.org/10.1109/ICDM.2016.0055 (2016).

16. Valstar, M. F. *et al.* Fera2017-addressing head pose in the third facial expression recognition and analysis challenge, DOI: https://doi.org/10.1109/FG.2017.107 (2017).

17. McCleery, A. *et al.* Meta-analysis of face processing event-related potentials in schizophrenia, DOI: https://doi.org/10.1016/j.biopsych.2014.04.015 (2015).

18. Graziano, P. A. & Garcia, A. Attention-deficit hyperactivity disorder and children's emotion dysregulation: A meta-analysis, DOI: https://doi.org/10.1016/j.cpr.2016.04.011 (2016).

19. Dalili, M., Penton-Voak, I., Harmer, C. & Munafò, M. Meta-analysis of emotion recognition deficits in major depressive disorder, DOI: https://doi.org/10.1017/S0033291714002591 (2015).

20. Poria, S. *et al.* Context-dependent sentiment analysis in user-generated videos, DOI: https://doi.org/10.18653/v1/P17-1081 (2017).

21. Baltrušaitis, T., Robinson, P. & Morency, L.-P. Openface: an open source facial behavior analysis toolkit, DOI: https://doi.org/10.1109/WACV.2016.7477553 (2016).

22. Sanh, V., Debut, L., Chaumond, J. & Wolf, T. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter, DOI: https://arxiv.org/abs/1910.01108 (2019).

23. Dham, S., Sharma, A. & Dhall, A. Depression scale recognition from audio, visual and text analysis, DOI: https://arxiv.org/abs/1709.05865 (2017).

24. Ngiam, J. *et al.* Multimodal deep learning, DOI: https://doi.org/10.5555/3104482.3104569 (2011).