PRONUNCIATION ASSESMENT

POWERED BY CTLC

OUR TEAM

Devansh Upadhyaya

DSEB

Vardan Vij

CSAI

Yashvi Maheshwari

DSEB

OUR PROBLEM?

This innovative machine learning model is set to transform English language learning by providing detailed feedback on fluency, accuracy, and pronunciation. It helps users identify areas for improvement and develop better communication skills, enhancing their mastery of the language.

OUR MOTIVATION

Bridging the Communication Gap

Aim to close the feedback gap in language learning, enhancing learners' English proficiency for better communication.

Advancing Educational Technology

Pushing ed-tech boundaries by using ML to personalize language learning, offering a scalable, adaptive learning tool.

Supporting Academic Excellence

Complementing CTLC's curriculum with a tool that boosts students' communication skills, preparing them for future success.

POTENTIAL APPLICATIONS

Customised Language Learning Apps
Integrate the model into mobile or web-based language learning applications to provide users with instant feedback on their speech, tailoring the learning experience to individual needs and accelerating progress in pronunciation and fluency.

Educational Institutions and Online Courses
Schools, universities, and online learning platforms can incorporate the model into their English language courses, offering students an additional resource for practicing and improving their spoken English skills outside of the traditional classroom setting.

Speech Therapy and Accent Reduction
The model could be adapted for use in speech therapy settings or accent reduction programs, assisting individuals in modifying specific pronunciation patterns or improving speech clarity and fluency for both therapeutic and professional development purposes.

POTENTIAL IMPACT

Enhanced Language Proficiency

The model could significantly improve learners' pronunciation, accuracy, and fluency in English, leading to better communication skills, greater confidence in speaking, and increased opportunities in global academic and professional arenas.

Innovative Educational Methods

By integrating advanced machine learning into language learning, the model paves the way for more personalised and efficient teaching methods, potentially revolutionising how languages are taught and learned across various educational platforms.

Cultural and Social Integration

Improved language skills facilitated by the model can ease the integration process for non-native speakers, fostering greater cultural exchange, understanding, and collaboration in increasingly multicultural and international communities.

LITERATURE SURVEY

SPEAKER FLUENCY LEVEL CLASSIFICATION USING ML

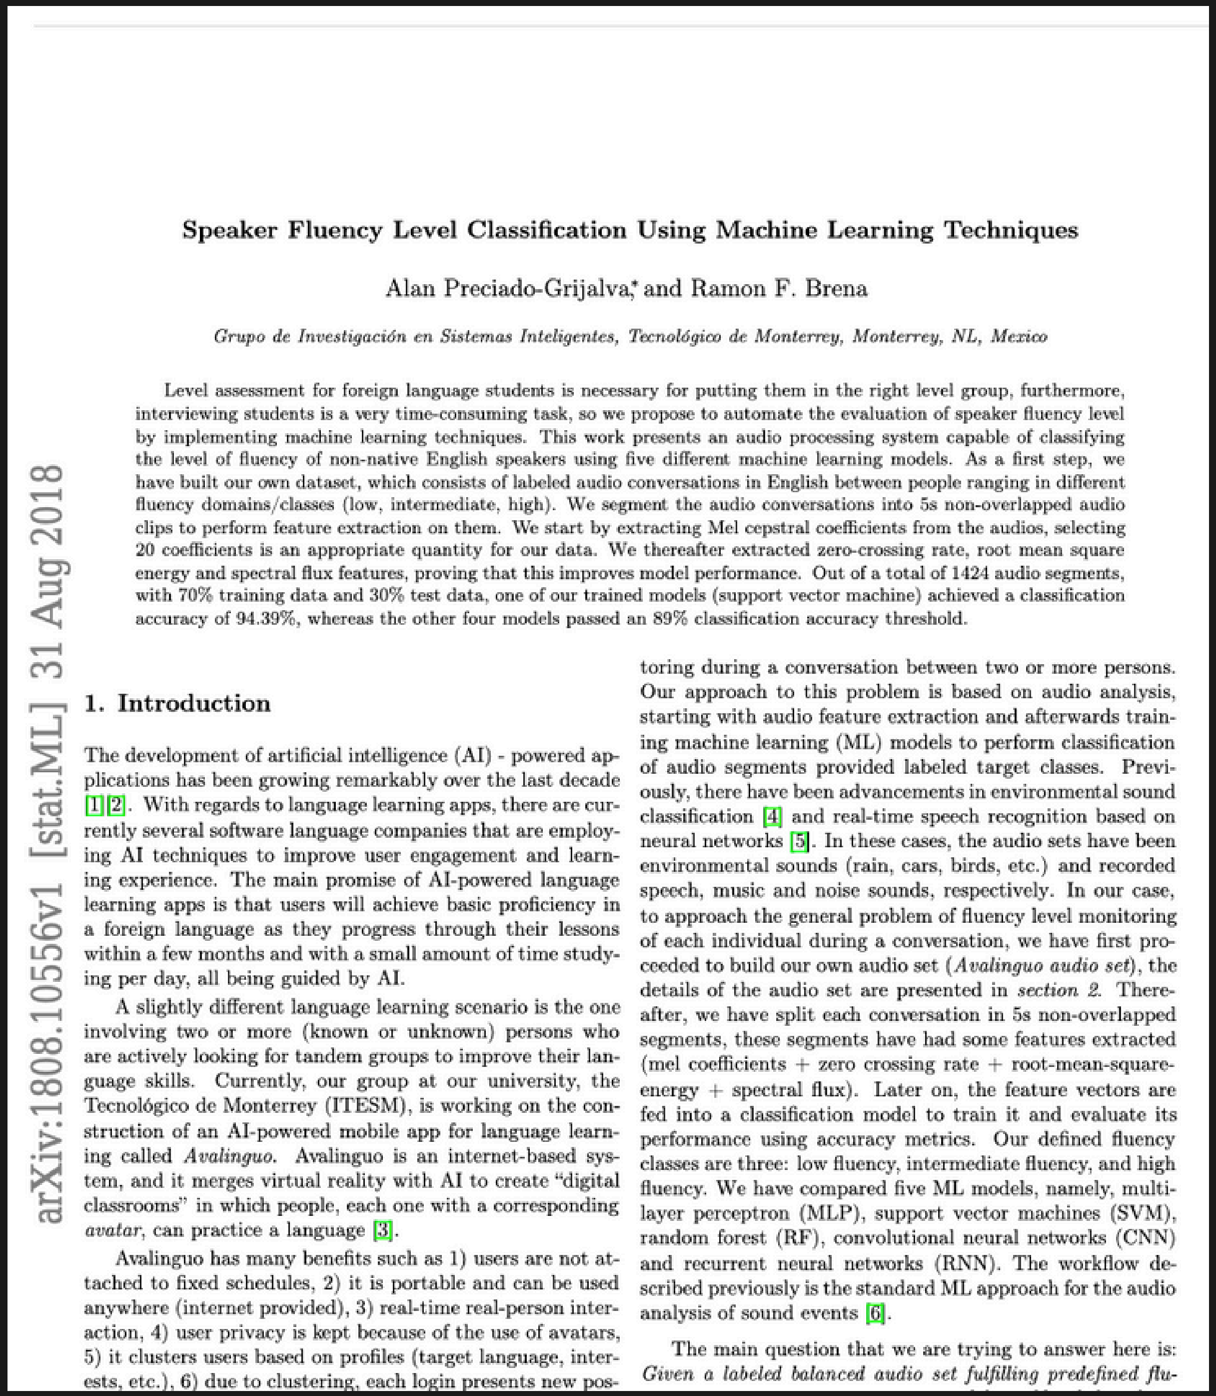

Speaker Fluency Level Classification Using Machine Learning Techniques

Alan Preciado-Grijalva,* and Ramon F. Brena

Grupo de Investigación en Sistemas Inteligentes, Tecnológico de Monterrey, Monterrey, NL, Mexico

Level assessment for foreign language students is necessary for putting them in the right level group, furthermore, interviewing students is a very time-consuming task, so we propose to automate the evaluation of speaker fluency level by implementing machine learning techniques. This work presents an audio processing system capable of classifying the level of fluency of non-native English speakers using five different machine learning models. As a first step, we have built our own dataset, which consists of labeled audio conversations in English between people ranging in different fluency domains/classes (low, intermediate, high). We segment the audio conversations into 5s non-overlapped audio clips to perform feature extraction on them. We start by extracting Mel cepstral coefficients from the audios, selecting 20 coefficients is an appropriate quantity for our data. We thereafter extracted zero-crossing rate, root mean square energy and spectral flux features, proving that this improves model performance. Out of a total of 1424 audio segments, with 70% training data and 30% test data, one of our trained models (support vector machine) achieved a classification accuracy of 94.39%, whereas the other four models passed an 89% classification accuracy threshold.

1. Introduction

The development of artificial intelligence (AI) - powered applications has been growing remarkably over the last decade [1][2]. With regards to language learning apps, there are currently several software language companies that are employing AI techniques to improve user engagement and learning experience. The main promise of AI-powered language learning apps is that users will achieve basic proficiency in a foreign language as they progress through their lessons within a few months and with a small amount of time studying per day, all being guided by AI.

A slightly different language learning scenario is the one involving two or more (known or unknown) persons who are actively looking for tandem groups to improve their language skills. Currently, our group at our university, the Tecnológico de Monterrey (ITESM), is working on the construction of an AI-powered mobile app for language learning called *Avalinguo*. Avalinguo is an internet-based system, and it merges virtual reality with AI to create "digital classrooms" in which people, each one with a corresponding *avatar*, can practice a language [3].

Avalinguo has many benefits such as 1) users are not attached to fixed schedules, 2) it is portable and can be used anywhere (internet provided), 3) real-time real-person interaction, 4) user privacy is kept because of the use of avatars, 5) it clusters users based on profiles (target language, interests, etc.), 6) due to clustering, each login presents new pos-

toring during a conversation between two or more persons. Our approach to this problem is based on audio analysis, starting with audio feature extraction and afterwards training machine learning (ML) models to perform classification of audio segments provided labeled target classes. Previously, there have been advancements in environmental sound classification [4] and real-time speech recognition based on neural networks [5]. In these cases, the audio sets have been environmental sounds (rain, cars, birds, etc.) and recorded speech, music and noise sounds, respectively. In our case, to approach the general problem of fluency level monitoring of each individual during a conversation, we have first proceeded to build our own audio set (*Avalinguo audio set*), the details of the audio set are presented in *section 2*. Thereafter, we have split each conversation in 5s non-overlapped segments, these segments have had some features extracted (mel coefficients + zero crossing rate + root-mean-square-energy + spectral flux). Later on, the feature vectors are fed into a classification model to train it and evaluate its performance using accuracy metrics. Our defined fluency classes are three: low fluency, intermediate fluency, and high fluency. We have compared five ML models, namely, multi-layer perceptron (MLP), support vector machines (SVM), random forest (RF), convolutional neural networks (CNN) and recurrent neural networks (RNN). The workflow described previously is the standard ML approach for the audio analysis of sound events [6].

The main question that we are trying to answer here is: *Given a labeled balanced audio set fulfilling predefined flu-*

Aim of Study

The primary aim of this study is to automate the process of evaluating non-native English speakers' fluency levels. By doing so, the study seeks to facilitate the placement of learners in appropriate level groups without the need for time-consuming personal interviews.

Dataset Used

The dataset created for this study, named the Avalinguo audio set, consists of labeled audio conversations in English, segmented into 5-second clips. These conversations cover a range of fluency levels (low, intermediate, high) and were designed to capture the variety of proficiency found among language learners

Model Used

Among the five machine learning models compared—Multi-layer Perceptron (MLP), Support Vector Machine (SVM), Random Forest (RF), Convolutional Neural Networks (CNN), and Recurrent Neural Networks (RNN)—the Support Vector Machine (SVM) was highlighted for its superior performance, achieving the highest classification accuracy of 94.39%.

Key Findings

The addition of features such as zero-crossing rate, root mean square energy, and spectral flux, alongside Mel cepstral coefficients, was found to enhance the performance of the models.

Reference

AUDIO-BASED ML MODEL FOR PODCAST LANGUAGE IDENTIFICATION

Aim of Study

The aim of the study was to develop an advanced machine learning model to accurately identify languages in podcast audio, addressing the unique challenges podcasts present.

Dataset Used

The dataset consisted of a wide variety of podcast episodes to ensure diversity in language, style, and format.

Model Used

The model used speaker embeddings in a two-step process for language identification, focusing on extracting relevant audio features that signify language characteristics.

Key Findings

The key finding was the model's high accuracy, with an average F1 score of 91.23%, indicating its effectiveness in identifying languages across different podcasts.

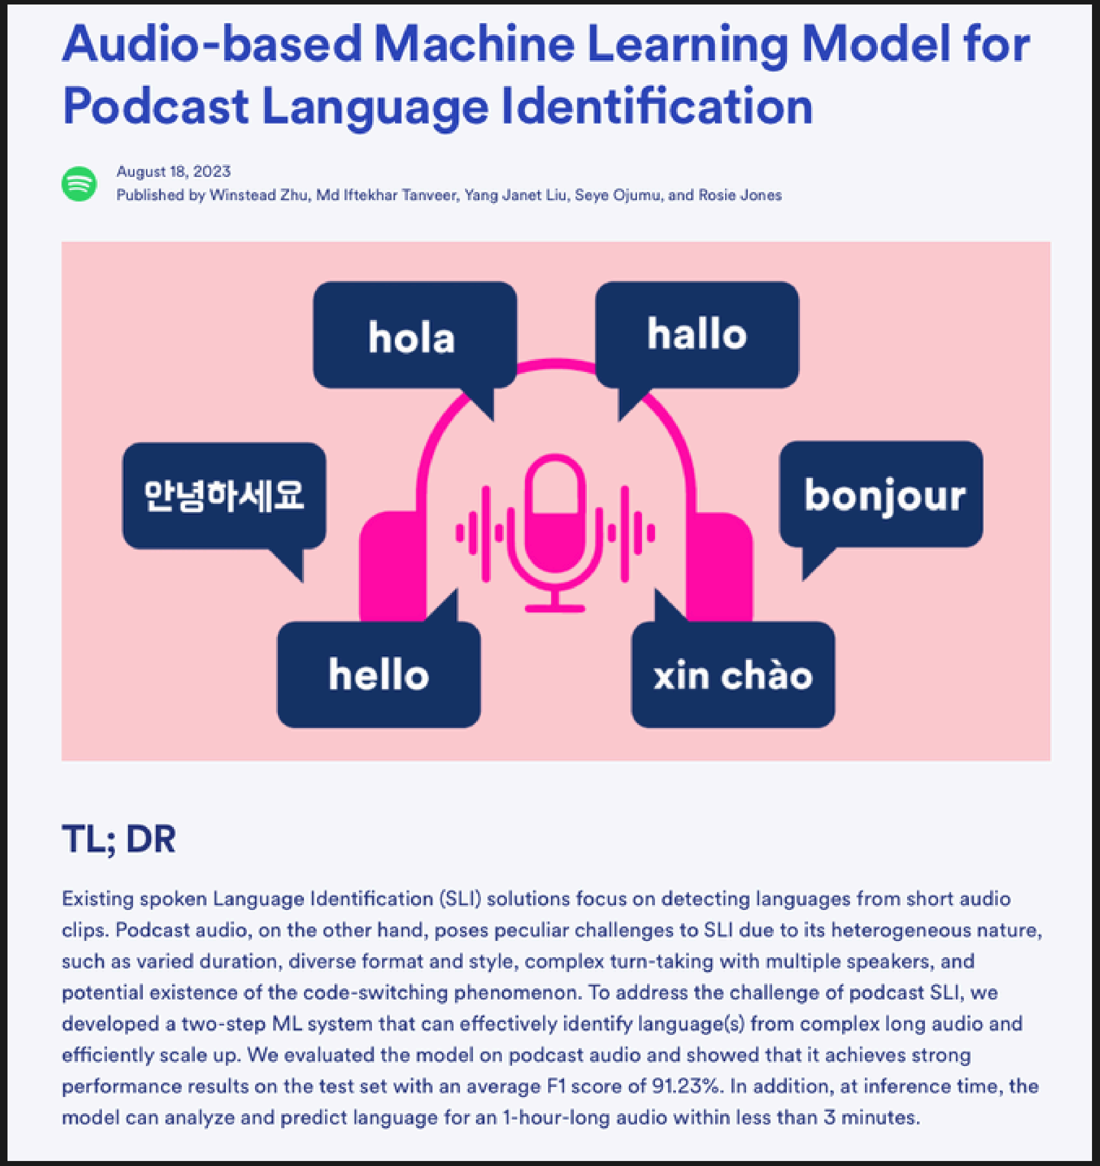

Audio-based Machine Learning Model for Podcast Language Identification

August 18, 2023
Published by Winstead Zhu, Md Iftekhar Tanveer, Yang Janet Liu, Seye Ojumu, and Rosie Jones

TL; DR

Existing spoken Language Identification (SLI) solutions focus on detecting languages from short audio clips. Podcast audio, on the other hand, poses peculiar challenges to SLI due to its heterogeneous nature, such as varied duration, diverse format and style, complex turn-taking with multiple speakers, and potential existence of the code-switching phenomenon. To address the challenge of podcast SLI, we developed a two-step ML system that can effectively identify language(s) from complex long audio and efficiently scale up. We evaluated the model on podcast audio and showed that it achieves strong performance results on the test set with an average F1 score of 91.23%. In addition, at inference time, the model can analyze and predict language for an 1-hour-long audio within less than 3 minutes.

Reference

DEEP LEARNING APPROACH FOR AN INDIAN ENGLISH REPOSITORY

Journal of Theoretical and Applied Information Technology
15th February 2024. Vol.102. No 3

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

AN END-TO-END DEEP LEARNING APPROACH FOR AN INDIAN ENGLISH REPOSITORY

PRATHIBHA SUDHAKARAN[1], ASHWANI KUMAR YADAV[2], SUNIL KARAMCHANDANI[3]

[1]Assisstant Professor, Muthoot Institute Of Technology and Science,Kochi,Kerala , India
[2]Assistant Professor , ASET,Amity University, Jaipur, Rajasthan, India
[3] Associate Professor, DJ Sanghvi College of Engineering, Mumbai, Maharashtra, India

E-mails: [1]prathibhasudhakaran@mgits.ac.in, [2]ashwaniy2@gmail.com, [3]skaramchandani@rediffmail.com

ABSTRACT

A voice recognition system that immediately translates raw audio waveforms into text without the need for separate components for language and acoustic modelling or manually constructed feature engineering is known as end-to-end deep learning for continuous speech recognition. It learns the whole audio to text mapping using a single deep neural network model. Conventional speech systems rely on intricate processing pipelines, but this method is far simpler. An end-to-end model in voice recognition is a simple single model that operates directly on words, subwords, or characters and may be trained from the ground up. This simplifies decoding by doing away with the requirement for both explicit phone modelling and a pronunciation lexicon. Deepspeech was used to construct and test the model, which was designed for Indian English. Additionally, a comparison is made between the results of the bi-directional RNN-based system and the traditional HMM model. With our method, we can quickly get a large amount of heterogeneous data for training due to a number of unique data synthesis techniques and an extremely effective RNN development system that utilises several GPUs. The connectionist temporal classification (CTC) objective function is used to infer the alignments between speech and label sequences, obviating the requirement for pre-generated frame labels. Experiments demonstrate that the RNN-based model has been observed to have equal word error rates (WERs) while also significantly speeding up the decoding process when compared to traditional hybrid HMM based on Kaldi.
Keywords: *End-To-End,HMM,CTC,RNN,CSR, Indian English, Deep Speech*

1 INTRODUCTION

The tasks involved in the field of speech recognition is difficult and complex. Teaching machines to comprehend and interpret spoken language is a huge technological achievement, despite the ease with which humans can do so. Speech variability, background noise, vocabulary and language models, contextual understanding, dependent and independent speaker models, and data availability and quality are some of the aspects that make speech recognition difficult. Despite these obstacles, recent developments in machine learning, deep learning, and neural network topologies have allowed for a considerable breakthrough in voice recognition technology. While attention approaches allow the model to focus on salient regions of the input signal, transformers facilitate the description of long-range relationships within the signal. These developments have greatly improved speech processing systems' functionality and performance, creating new prospects for use in a variety of industries.[1]

Despite significant advancements in speech processing, deep learning still has some issues that need to be resolved. These difficulties include the need for large amounts of labeled data, the models' interpretability, and their resilience to various
environmental factors. Deep learning approaches have significantly improved the analysis, synthesis, and detection of speech signals, which are all parts of speech processing. Many researchers previous workhave shed light on the

Aim of Study

The study aimed to develop an end-to-end deep learning model for continuous speech recognition that translates raw audio waveforms directly into text without requiring separate components for language and acoustic modeling or manual feature engineering.

Dataset Used

The dataset for this study was provided by the Indian Institute of Technology Madras (IITM) and included Indian English spoken by male and female speakers. The dataset comprised 45,000 tokens.

Model Used

The model employed was based on a recurrent neural network (RNN) trained to process voice spectrograms and output text transcriptions. The RNN model consisted of five dense layers and one unidirectional RNN layer, utilizing MFCC (Mel-Frequency Cepstral Coefficients) features extracted from the audio samples.

Key Findings

The study found that the RNN-based model could achieve comparable word error rates (WERs) while significantly speeding up the decoding process.

Reference

CURRENT IMPLICATIONS

Technology Utilised
Speech Recognition, Natural Language Processing
Application
Duolingo uses speech recognition technology to assess users' pronunciation by comparing it to expected phonetic patterns. It provides immediate feedback to help learners correct and improve their spoken language skills.

Technology Utilised
Proprietary Speech Recognition Engine
Application
Rosetta Stone's TruAccent speech-recognition engine is designed to provide real-time pronunciation feedback, helping learners to speak confidently and accurately.

Technology Utilised
BERT (Bidirectional Encoder Representations from Transformers)
Application
Google has been exploring the use of BERT in various language understanding tasks, which can potentially be applied to pronunciation assessment to analyse context and provide more accurate feedback on language use.

Technology Utilized
Machine Learning, Deep Learning Algorithms
Application
ELSA Speak uses advanced speech recognition technology powered by deep learning algorithms to listen to the way language learners pronounce words, providing expert feedback on pronunciation errors and how to fix them.

DATASET AND PRE-PROCESSING

OUR DATA

- *We used the* ***Mozilla Common Voice*** *dataset from Kaggle for our project.*

- *We chose this dataset because it was the most extensive, covering a broad range of everyday pronunciations and phrases. Its exceptional diversity made it the best fit for training our model to understand various accents and linguistic nuances.*

- *Compared to the LibriSpeech and CMU Pronunciation datasets, Mozilla Common Voice offered a richer variety of real-world speech patterns. This included diverse accents and colloquial expressions, which were crucial for developing a versatile pronunciation assessment tool.*

- *The dataset is ethically sourced. Mozilla emphasizes volunteer contributions and explicit consent for use in projects. This ensures the respect for privacy and the fair use of participants' data.*

Targets	Justification
Fluency Score	*Evaluates rhythm and speed, indicating how naturally speech flows.*
Accuracy Score	*Measures correctness of language use, reflecting the speaker's grasp of English.*
Completeness Score	*Assesses if responses are fully formed, gauging comprehension and response adequacy.*
Pronunciation Score	*Rates the clarity of speech sounds, crucial for intelligibility and communication*

Features	Justification
Average Pitch	*Shows the speaker's tone, which can affect the perceived emotion and assertiveness of speech*
Average Intensity	*Measures loudness, important for understanding speech dynamics and energy.*
Jitter	*Indicates frequency stability, with higher values suggesting potential voice disorders or tension*
Shimmer	*Reflects amplitude stability; variations can signal vocal strain or emotional stress.*
MFCC	*Captures timbral aspects, essential for identifying unique speech characteristics and speaker differences.*

CHALLENGES FACED

1. Computational Power and Time
2. Imbalance in Dataset

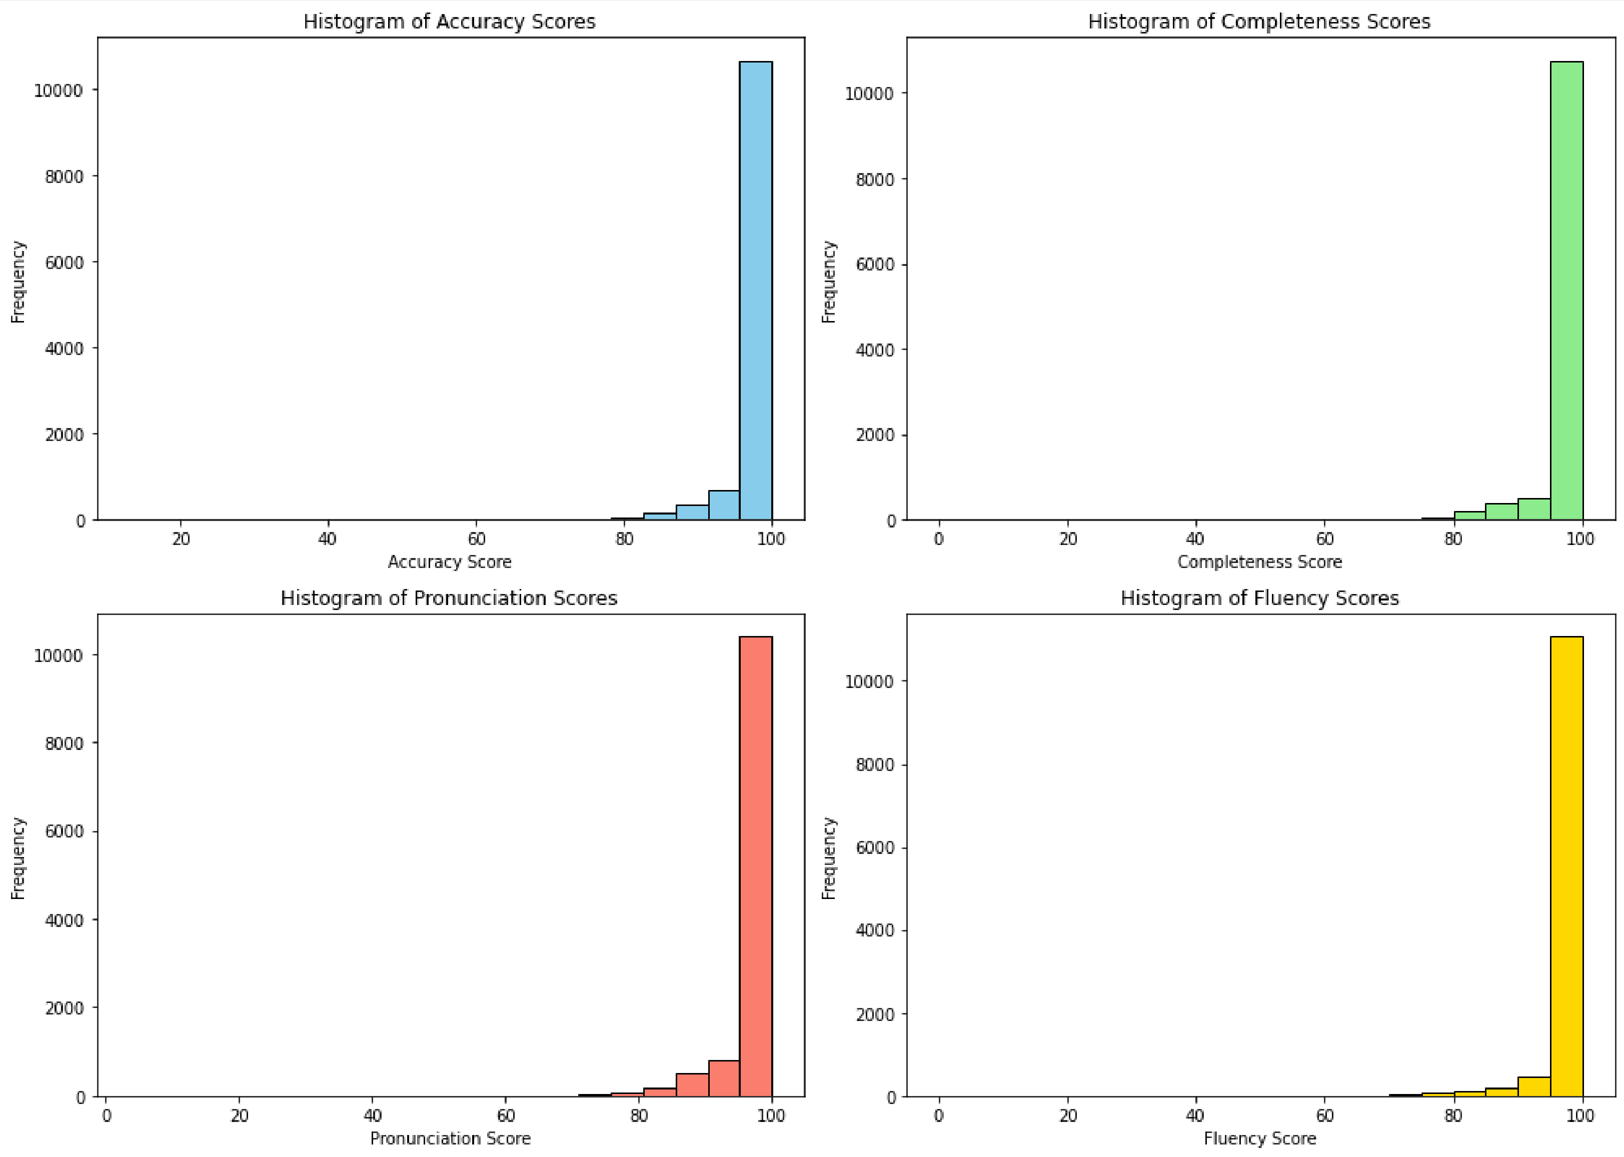

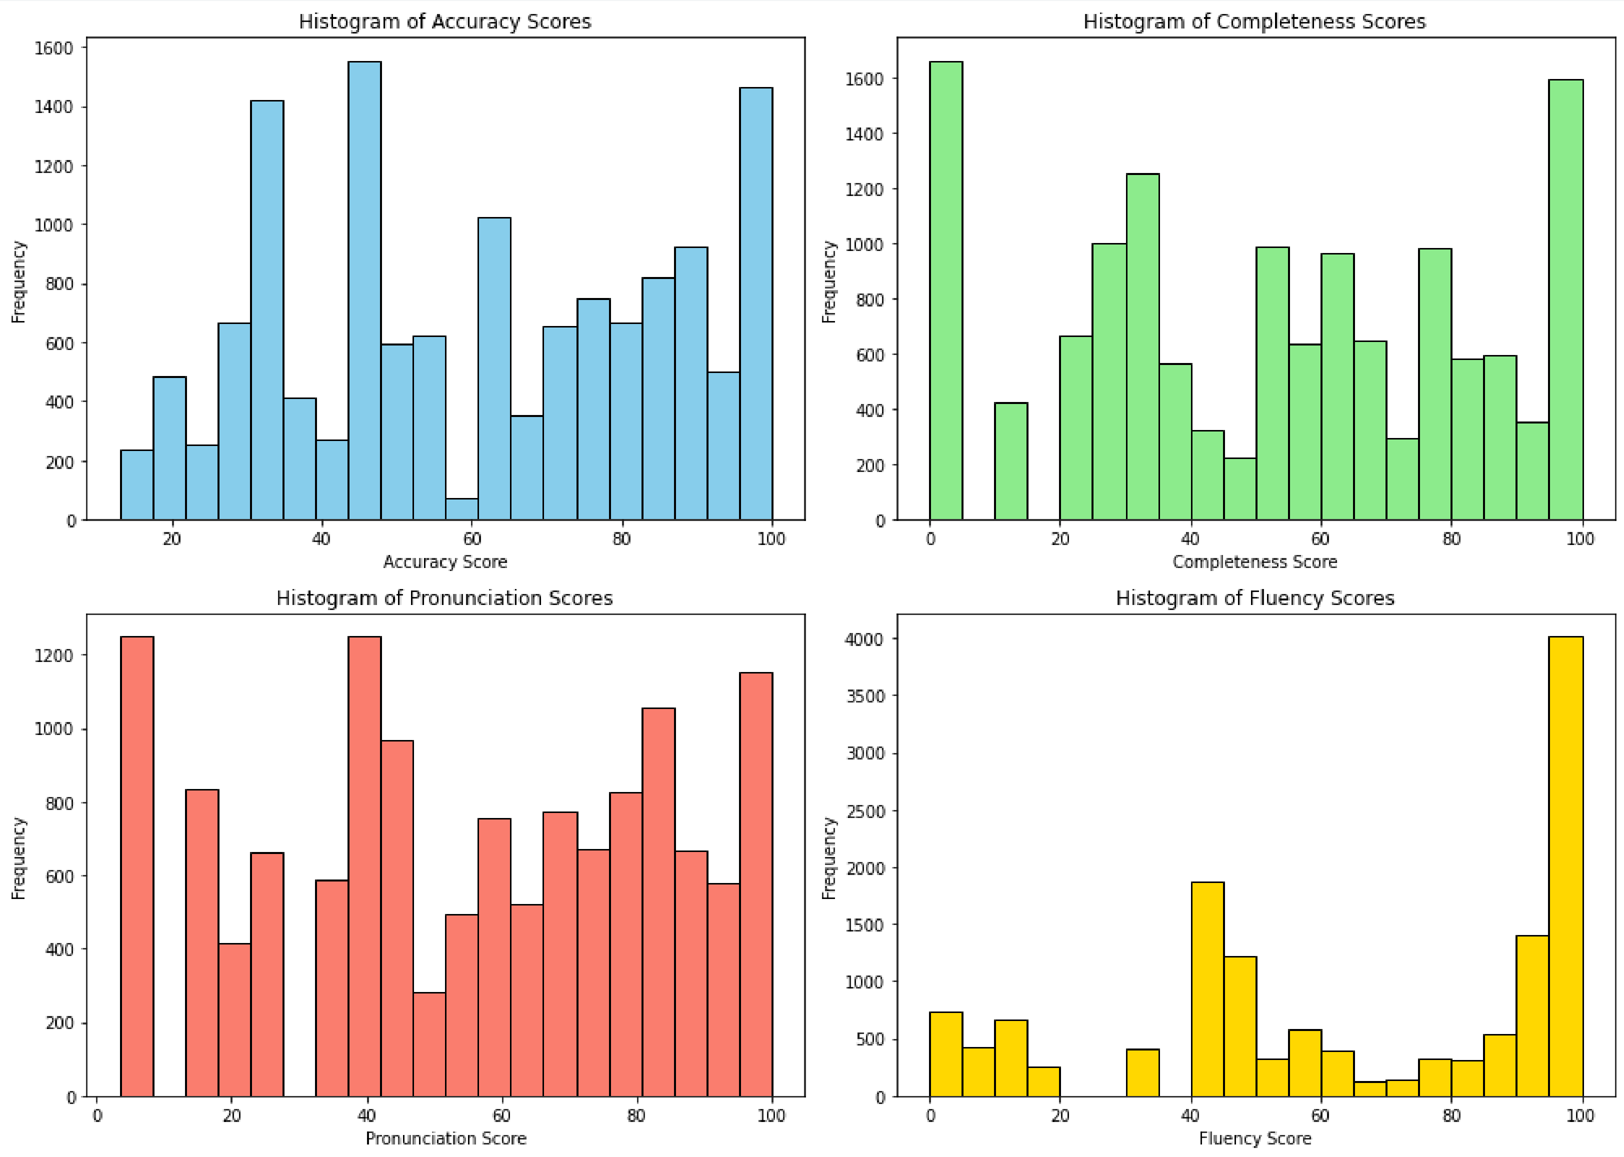

OUR DATA - RAW

	A	B
1	filename	text
2	sample-000000	learn to recognize omens and follow them the old king had said
3	sample-000001	everything in the universe evolved he said
4	sample	
5	sample	
6	sample	
7	sample	
8	sample	
9	sample	
10	sample	
11	sample	
12	sample	
13	sample	
14	sample	
15	sample	
16	sample	
17	sample	
18	sample	
19	sample	
20	sample	
21	sample	
22	sample	
23	sample	
24	sample	
25	sample	
26	sample	
27	sample	
28	sample	
29	sample	
30	sample	
31	sample	
32	sample	
33	sample	
34	sample	
35	sample	
36	sample	

195744	sample-195742	all that time the martians must have been getting ready
195745	sample-195743	nobody knew who malcolm was
195746	sample-195744	the boy went to look for the englishman
195747	sample-195745	all the joy he had seen that morning had suddenly disappeared
195748	sample-195746	but they really don't know what they're saying
195749	sample-195747	i can personally vouch for his character
195750	sample-195748	i gotta figure some way out of this thing
195751	sample-195749	so they think she left it with you
195752	sample-195750	in his youth my dad was a hooligan
195753	sample-195751	i need you to be spontaneous he asked me out to din din
195754	sample-195752	and stop threatening that boy
195755	sample-195753	two more months passed and the shelf brought many customers into the crystal shop
195756	sample-195754	the silicon sealant has dried
195757	sample-195755	he just got a new kite for his birthday
195758	sample-195756	then suddenly they kneel and die
195759	sample-195757	some of those who saw its flight say it travelled with a hissing sound
195760	sample-195758	ogilvy told him everything that he had seen
195761	sample-195759	the snake fought frantically making hissing sounds that shattered the silence of the desert
195762	sample-195760	tell them to start getting those extras out
195763	sample-195761	stand by to contact amanda
195764	sample-195762	will you get out and let me handle this
195765	sample-195763	it seemed as if what the old king had called beginner's luck were no longer functioning
195766	sample-195764	i'll work for you he said
195767	sample-195765	it was harassment pure and simple
195768	sample-195766	but before i go i want to tell you a little story
195769	sample-195767	down below in the darkness were hundreds of people sleeping in peace
195770	sample-195768	i have been waiting for you here at this oasis for a long time
195771	sample-195769	the burning fire had been extinguished
195772	sample-195770	he heard a muffled grating sound and saw the black mark jerk forward an inch or so
195773	sample-195771	the englishman said nothing
195774	sample-195772	the irish man sipped his tea
195775	sample-195773	what do you know about that
195776	sample-195774	the phone rang while she was awake
195777	sample-195775	among these people were a couple of cyclists a gardener i employed sometimes and a girl carrying a baby

OUR DATA - PROCESSED

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
filename	MFCC_1	MFCC_2	MFCC_3	MFCC_4	MFCC_5	MFCC_6	MFCC_7	MFCC_8	MFCC_9	MFCC_10	MFCC_11	MFCC_12	MFCC_13	MFCC_14	MFCC_15	MFCC_16	MFCC_17	MFCC_18	MFCC_19	MFCC_20
sample-000000	-267.0175	111.51408	-25.76993	22.131796	1.965262	-4.992992	0.2580115	-14.19481	-2.923832	-2.122391	-6.857883	2.1503375	-8.371711	-8.446826	-3.338415	-7.642871	-14.40806	-5.088026	-7.608096	0.1366174
sample-000001	-695.7019	73.624504	10.015617	15.413986	2.9955633	-2.319246	-2.690451	-0.899252	-6.909869	0.6489629	-3.922787	-0.784457	-3.011517	-0.728425	-4.652588	0.4625341	-5.079135	-1.769182	-4.434768	-1.631329
sample-000002	-330.9641	55.5857	20.686192	33.323223	13.275977	10.631326	-7.737192	-0.644539	-12.52795	-1.1271	-6.412432	-5.580572	-23.22936	3.3521557	4.807851	-4.014465	-5.240258	-2.495067	-5.085029	-2.235309
sample-000003	-526.975	92.90594	9.87802	12.6751	2.5505998	-0.936728	-10.34216	-13.46044	-8.697995	-6.849994	-2.936985	-4.727073	-4.78264	0.4819029	-5.765874	-3.141343	-4.779899	-0.729745	-7.412266	-2.991007
sample-000004	-72.50391	84.04445	-5.079901	21.510885	8.788865	4.4337215	-3.843855	0.3073304	-3.926708	-8.768681	-9.724998	-4.622541	-11.88935	-0.878991	-9.199817	-9.92968	-11.13104	-11.38618	-11.82231	-6.413286
sample-000005	-576.3187	86.37545	11.825168	49.36195	-5.830213	-0.714912	-14.10705	-8.644947	-2.919143	-14.07632	-14.48186	-13.57874	-6.720933	-6.186358	-4.839341	-9.007878	-3.113941	-1.743996	-9.447626	-8.065131
sample-000006	-292.0289	81.23502	16.444843	19.271765	15.7108	3.202271	4.964137	-1.770349	-3.299132	-0.971218	-5.853793	-3.439825	-7.764057	-7.646153	1.2282255	2.2814918	-7.068227	2.195734	-6.521924	1.7849219
sample-000007	-148.8642	130.72702	-8.807579	4.902917	29.032015	6.9800463	2.8942962	7.7400913	-14.15415	6.803039	-5.873151	4.4977603	-0.251702	5.515161	-6.32646	4.5125523	-2.759242	3.4155493	-7.939445	4.161499
sample-000008	-414.3074	97.79011	-25.67946	8.567528	9.028595	-0.727907	-20.75136	-23.03922	-14.79138	-9.397761	1.9272534	4.4504604	-5.883214	-4.832831	-12.1924	-2.329786	-6.104097	4.5657945	-5.780686	-2.0063
sample-000009	-269.8825	74.74891	-30.72167	91.42908	-14.67624	-17.17333	-5.024773	-15.65498	-4.521118	-9.562298	-11.59298	-7.057419	-1.486495	-7.617592	4.6963897	-7.390301	-1.658879	2.4062438	-11.4466	3.1970258
sample-000010	-416.3499	87.90993	-13.67994	28.17101	-7.17555	-2.995975	-6.759255	-4.127787	-2.535561	-10.50133	-6.277811	-11.0682	-7.567941	-4.096524	-6.666657	-6.413778	-7.401173	-12.14971	-11.83185	-11.02761

OUR DATA - LABELLED

	A	V	W	X	Y	Z	AA	AB	AC	AD
1	filename	accuracy_score	fluency_score	completeness_score	pronunciation_score	average_pitch	average_intensity	speech_rate	jitter	shimmer
2	sample-000000	100	99	100	99.4	111.7	63.3	175.4	2.7	2.2
3	sample-000001	99	76	100	85.4	140.5	20	55.6	5.9	4.4
4	sample-000002	100	100	100	100	153	58.9	165.9	5.2	2.5
5	sample-000003	100	92	100	95.2	223.4	40.2	135.4	3.9	3.8
6	sample-000004	100	100	100	100	213.3	81	208.3	2.7	1.4
7	sample-000005	100	95	100	97	197.5	38.4	144	2.1	4
8	sample-000006	95	91	100	93.6	119.5	61	115	2	2.3
9	sample-000007	100	99	100	99.4	127.7	73.4	140.8	4.4	2.6
10	sample-000008	100	100	100	100	181.5	50.3	208.3	8	3.5
11	sample-000009	96	96	100	96.8	146.1	65	148.4	2.1	1.8
12	sample-000010	100	84	100	90.4	212.4	49.6	115.2	1.8	3

VISUALISATIONS

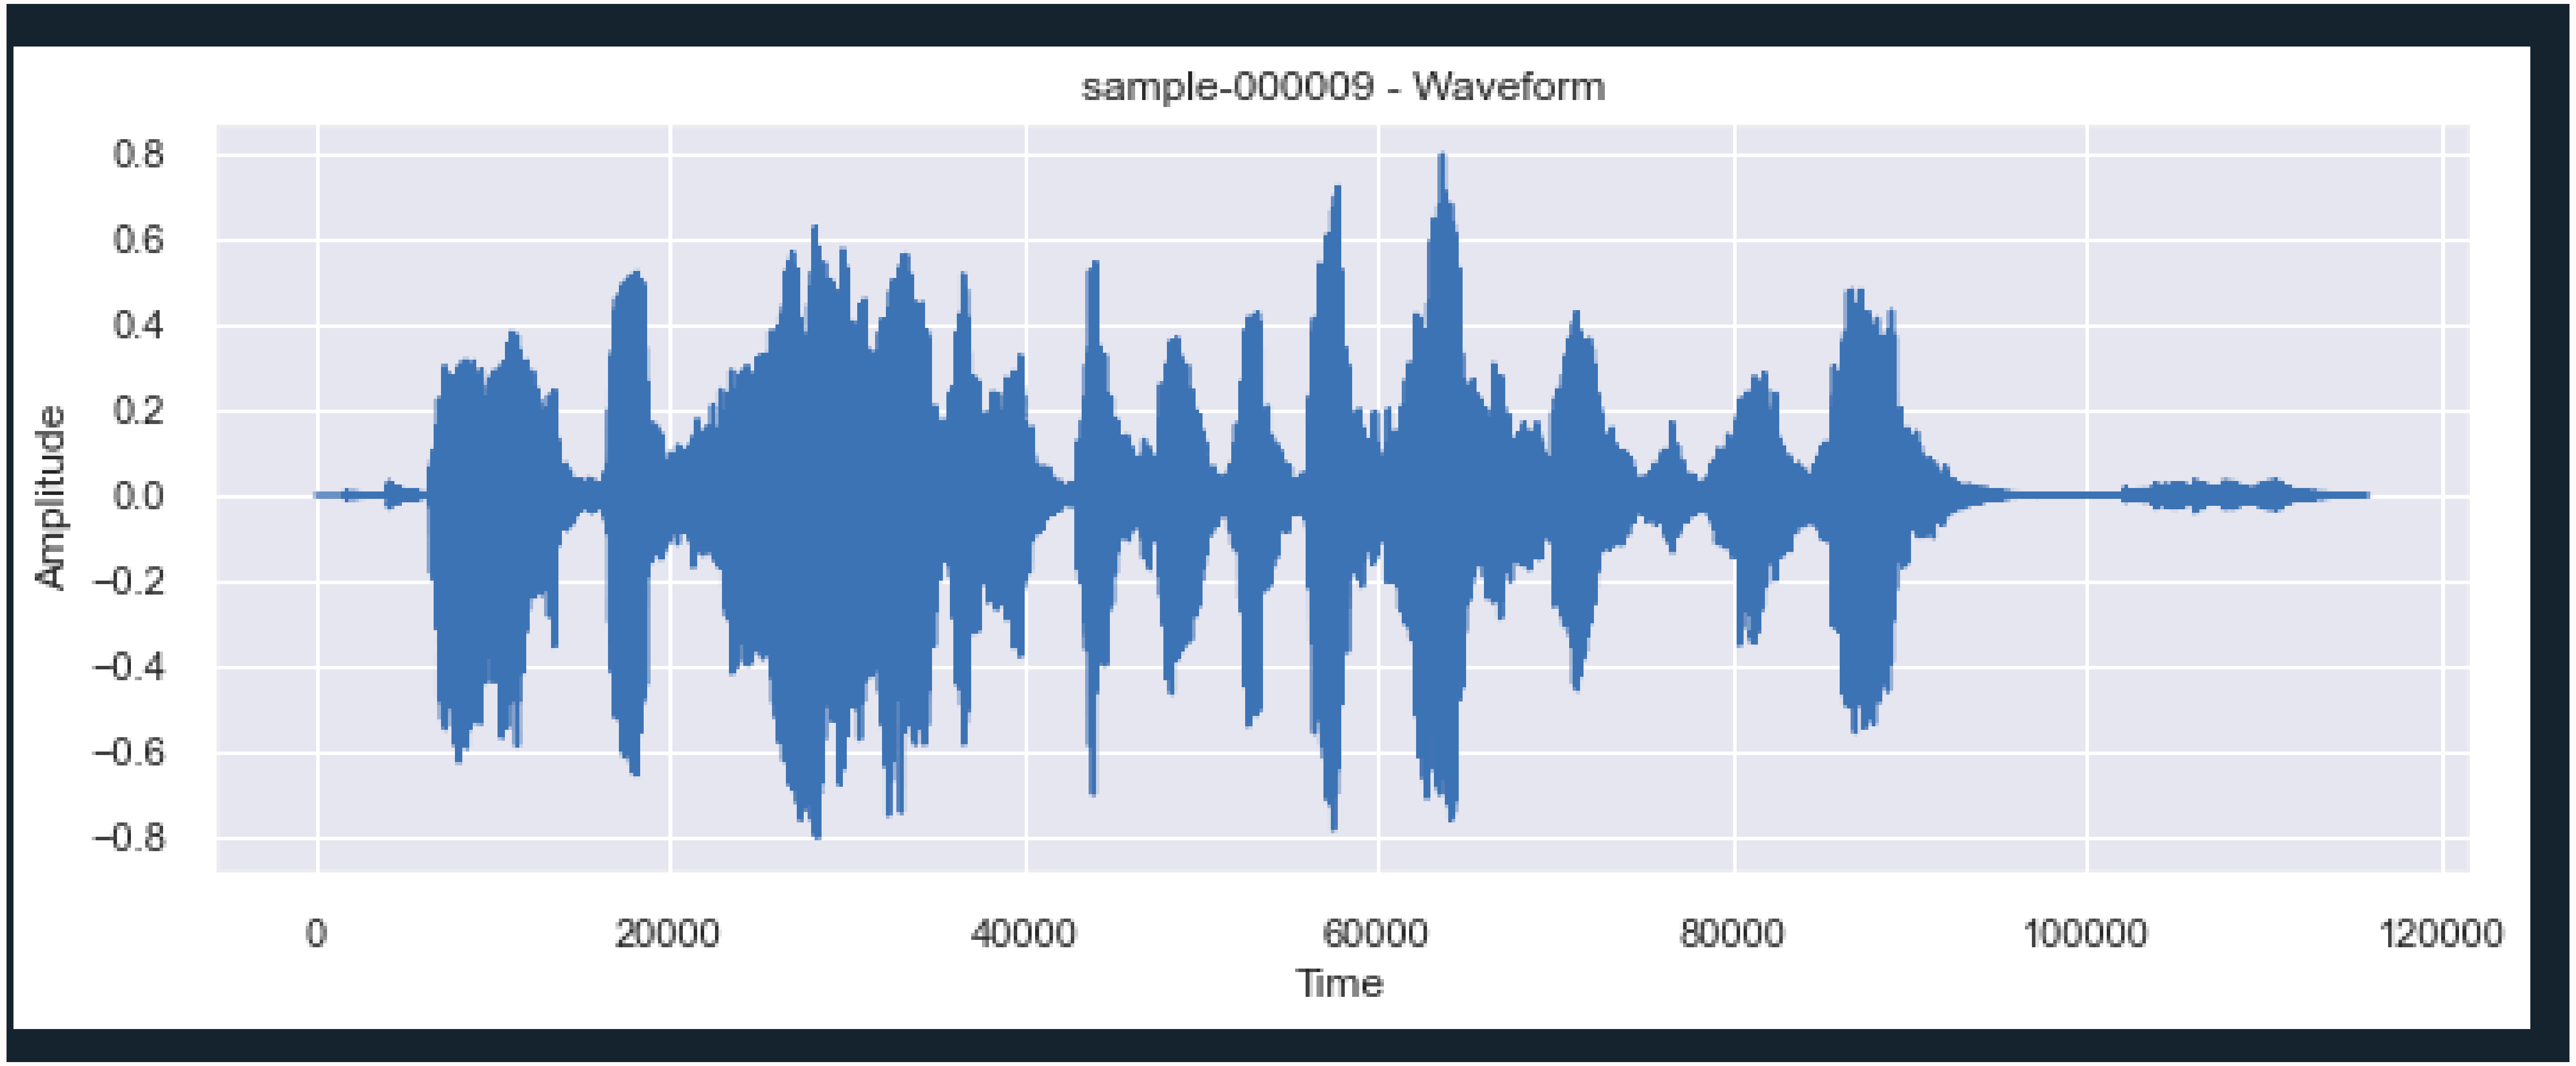

VISUALISATIONS

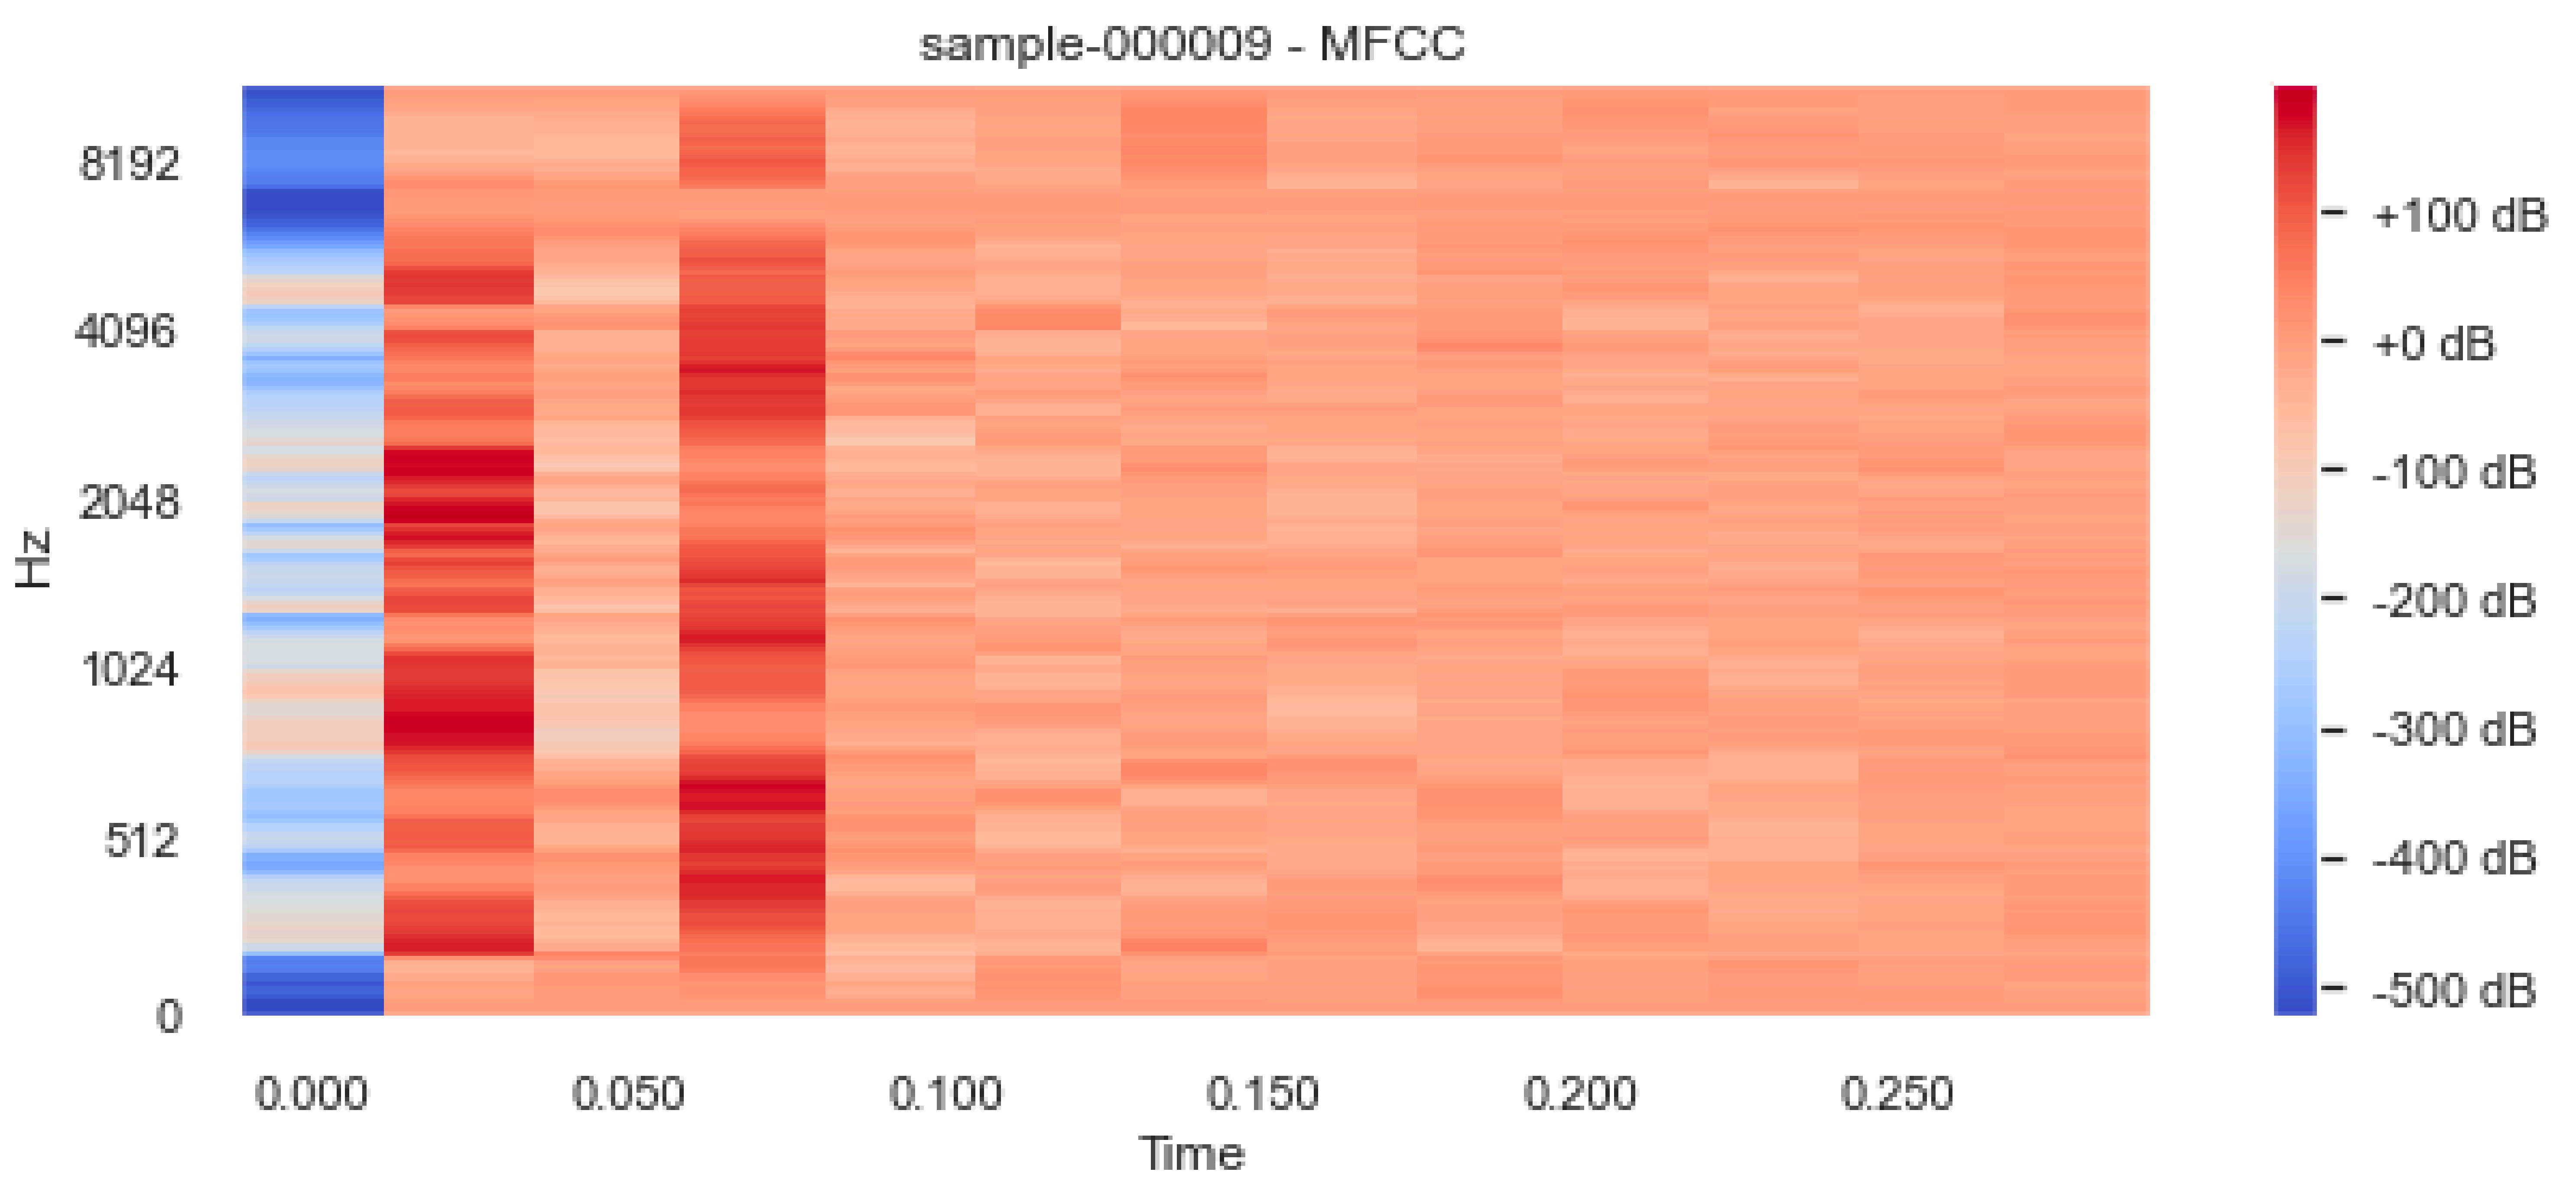

OUR MACHINE LEARNING MODEL

LSTM
LONG SHORT-TERM MEMORY MODEL

- *LSTMs excel at capturing temporal dependencies in sequential data like speech.*
- *Their ability to selectively remember or forget information makes them well-suited for modeling the dynamic nature of speech patterns.*
- *Unlike SVMs, which are primarily designed for fixed-length feature vectors, LSTMs can handle variable-length speech segments.*
- *While other neural networks lack the memory mechanisms crucial for understanding speech context,LSTMs are explicitly designed to retain relevant information over time.*

OUR MODEL

1.Data Processing
2. Train-Test Split (Train: 0.6, Test:0.2, Validation:0.2)
3. Input Layer
4. LSTM Layer: Processes the sequential input data. LSTM layers are effective for sequence prediction problems because they can maintain information in memory over time.
5. Dropout Layer: Helps in preventing overfitting by randomly setting a fraction of input units to 0 during training.
6. Output Layer: Outputs the predictions for each target score. This layer uses a linear activation function to allow for a continuous range of outputs.
7. Training
8. Evaluation
9. Testing

MODEL

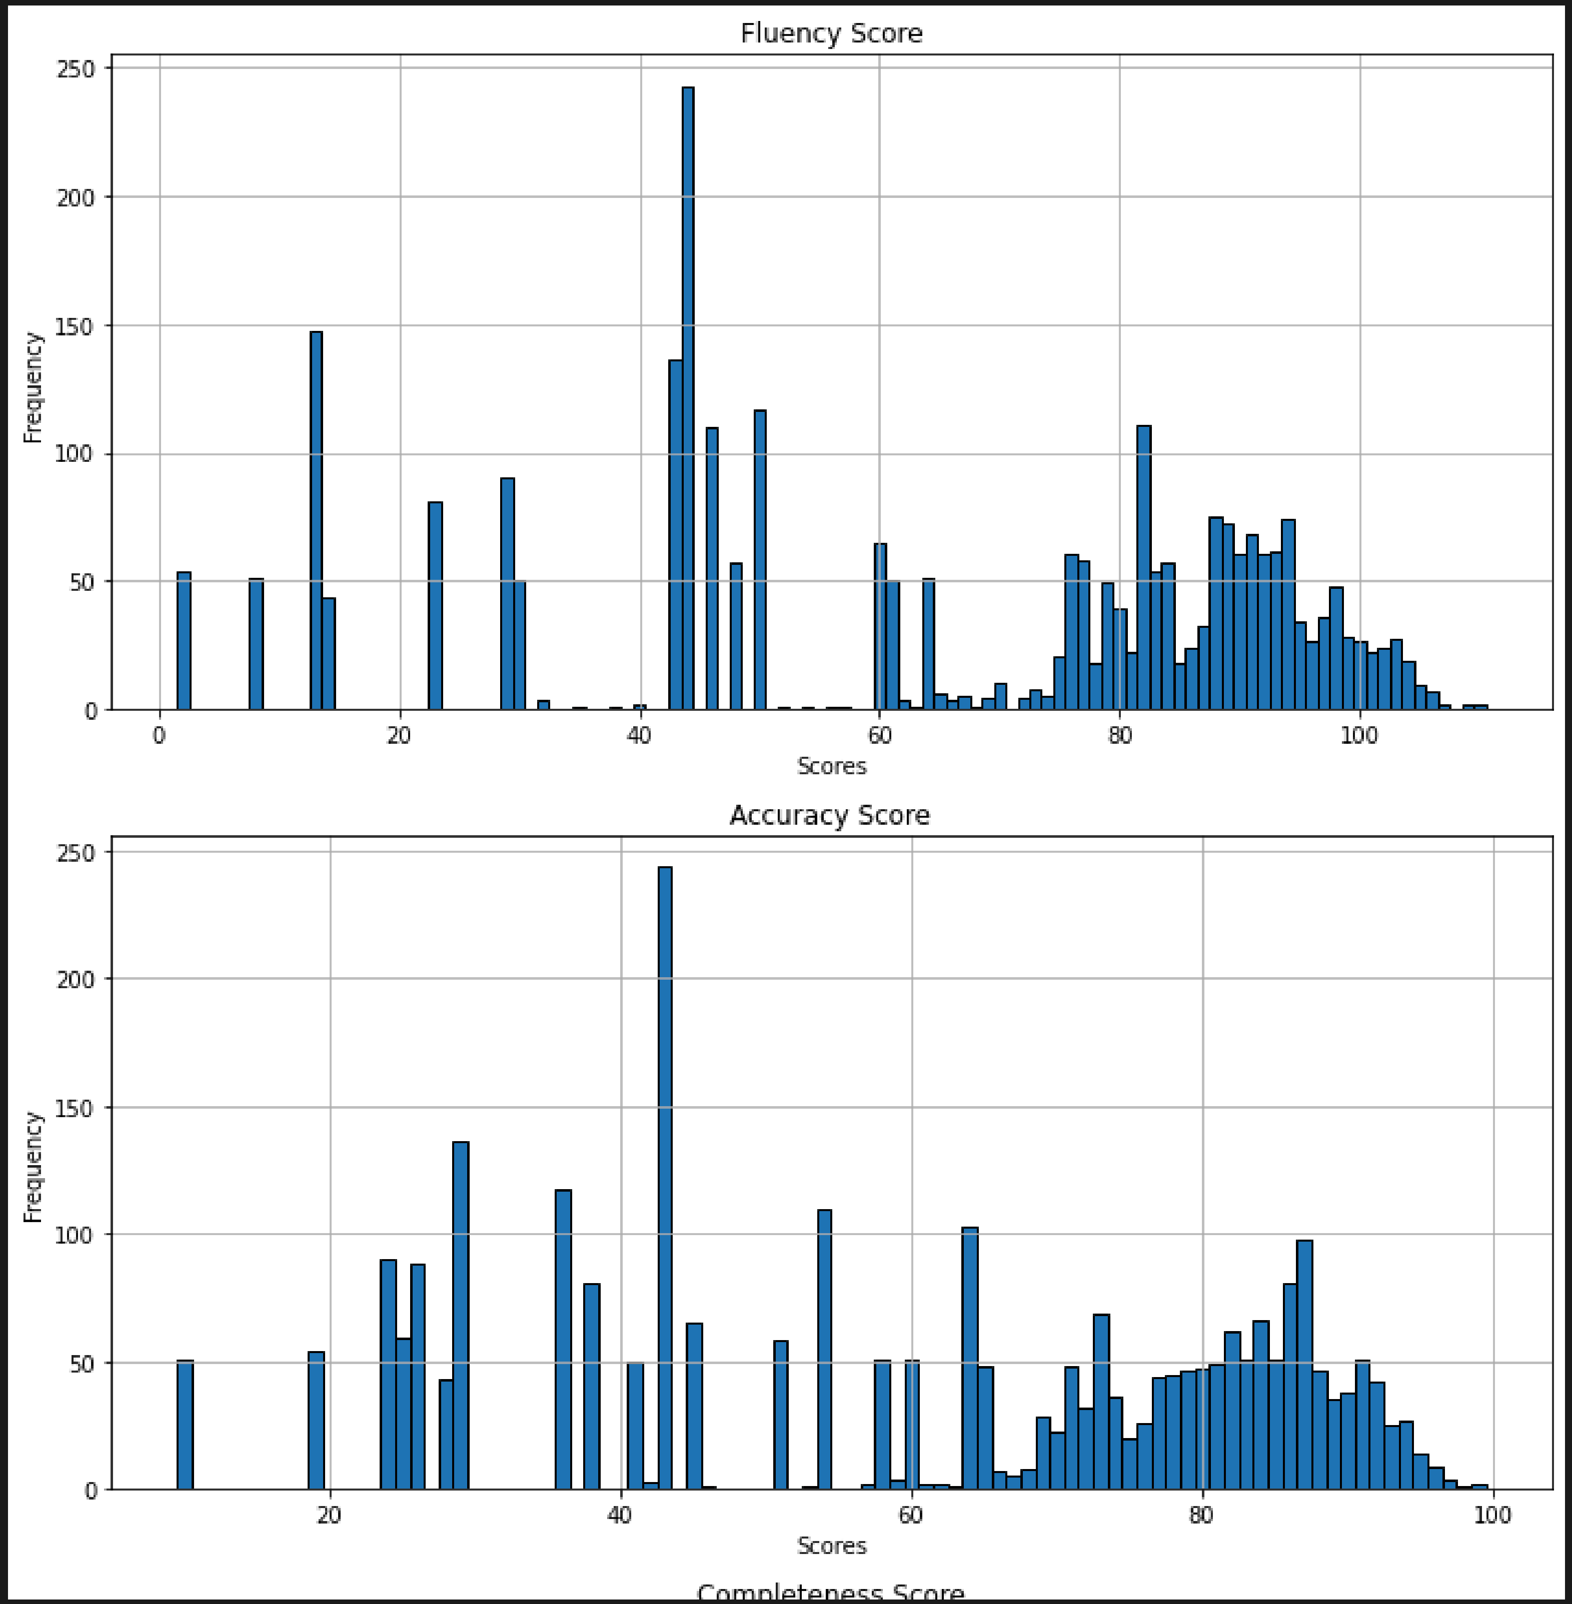
Fluency Score
Frequency
Scores
Accuracy Score
Frequency
Scores
Completeness Score

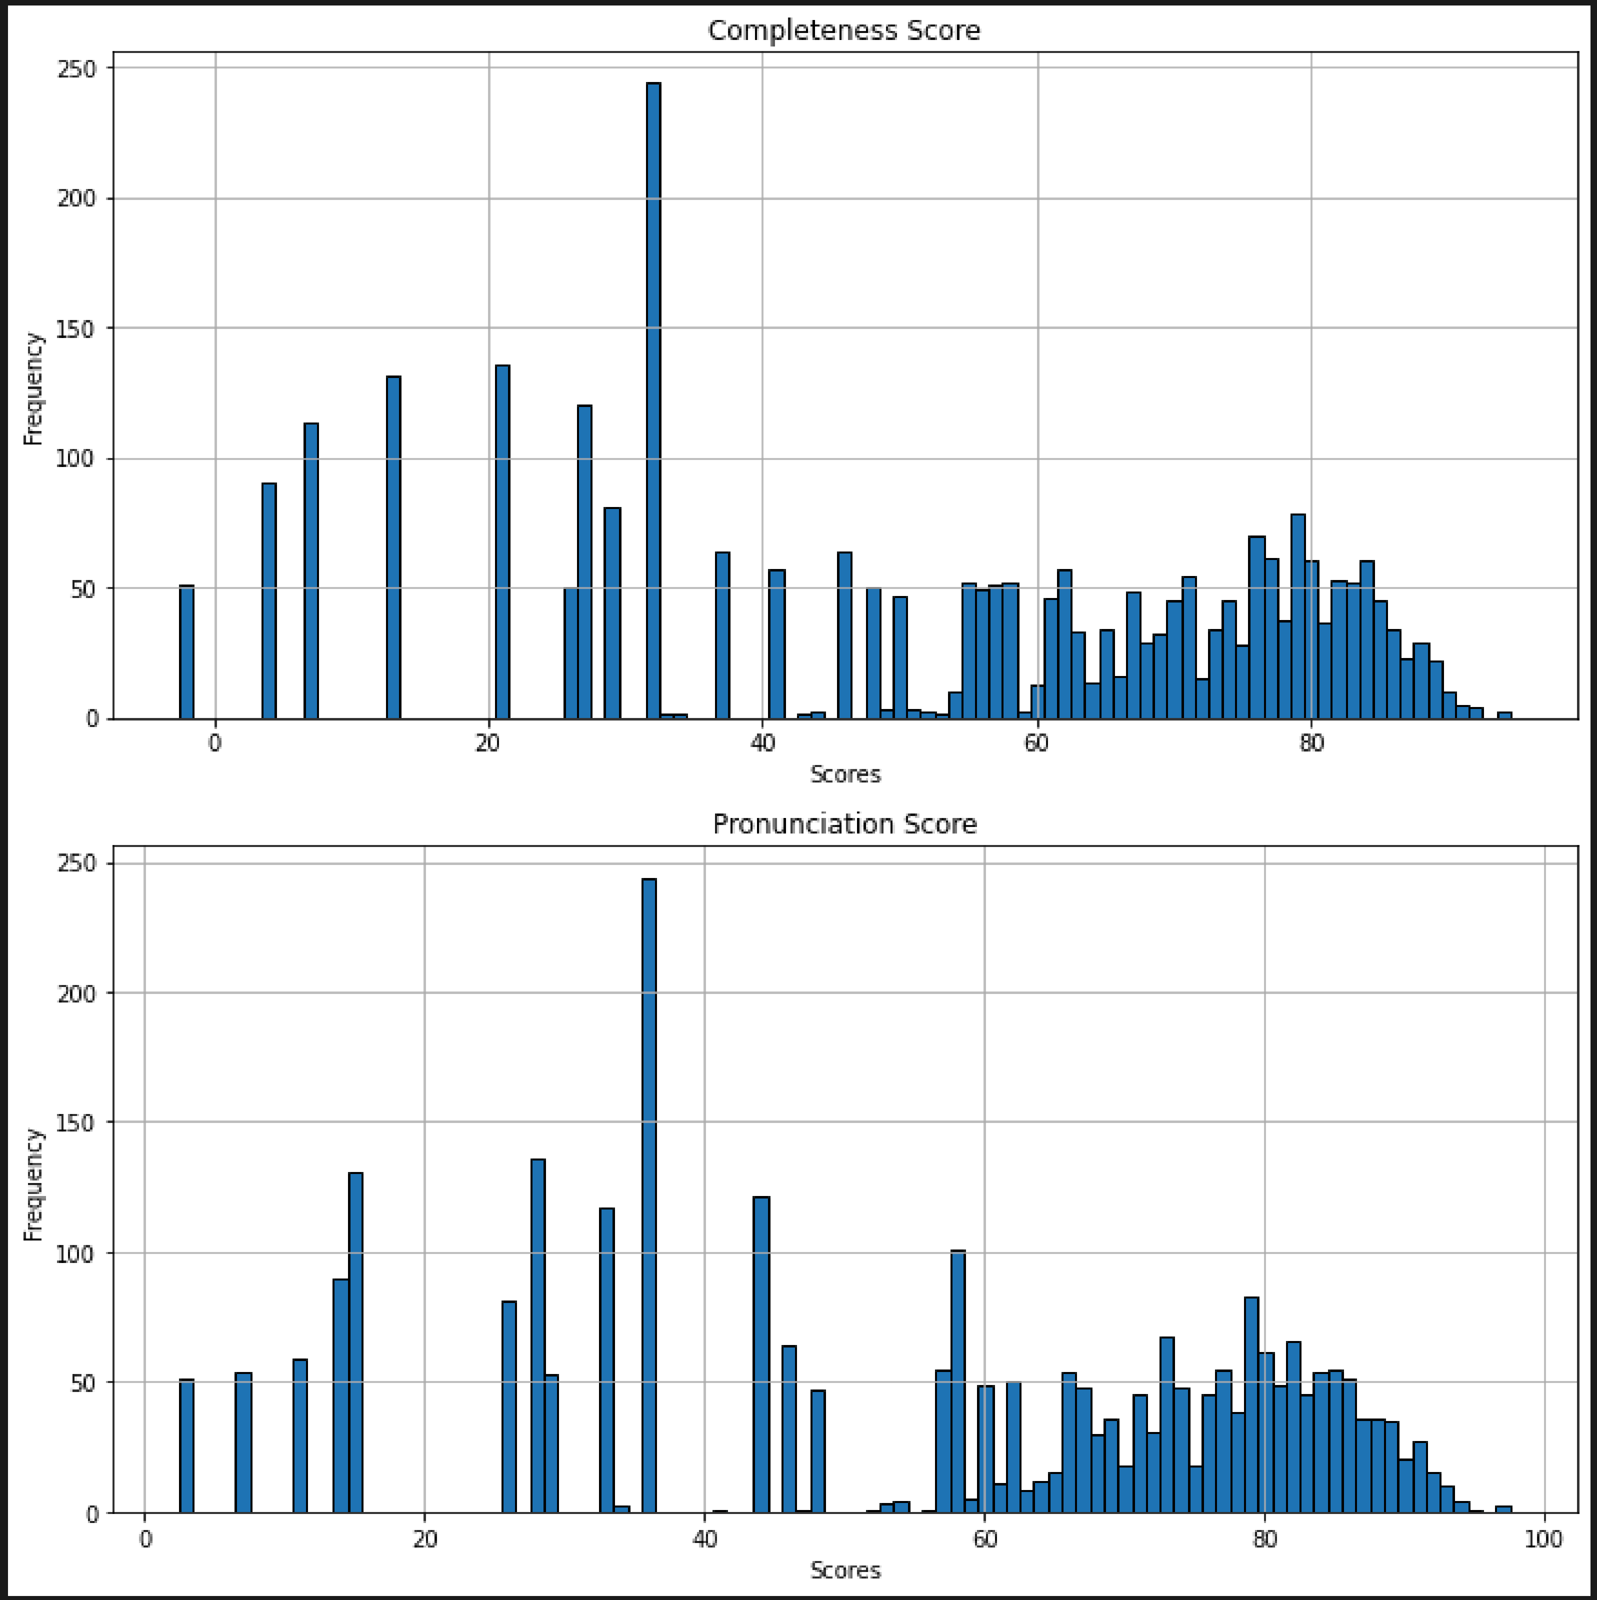
Completeness Score
Frequency
Scores
Pronunciation Score
Frequency
Scores

REAL WORLD TESTS

```
Sample 4 - Filename: sample-014475
Predictions:
  Fluency Score: 43.40
  Accuracy Score: 29.72
  Pronunciation Score: 21.71
  Completeness Score: 28.08
Actual Scores:
  Fluency Score: 47.0
  Accuracy Score: 27.0
  Pronunciation Score: 20.0
  Completeness Score: 26.8
```

```
Sample 5 - Filename: sample-002947
Predictions:
  Fluency Score: 48.12
  Accuracy Score: 51.05
  Pronunciation Score: 41.79
  Completeness Score: 44.47
Actual Scores:
  Fluency Score: 47.0
  Accuracy Score: 51.0
  Pronunciation Score: 42.0
  Completeness Score: 44.8
```

REAL WORLD TESTS

Predicted scores: [76.885994 60.845703 51.427834 57.933548]

IMPLEMENTATION

Our Project directly aligns with revolutionary styles of teaching at Plaksha University.

This model will be a great value addition at CTLC lab. Where students can perfect their pronunciations and also improve their speaking skills.

We envision this model to give you challenging phrases, level by level to which can help improve Vocabulary skills.

CHALLENGES

Computation

Generalisation

UI/UX

PERFORMANCE METRICS

Test Loss: 109.37371826171875

Test MAE: 7.058261394500732

THANK YOU!