

End-to-end Triplet Loss based Emotion Embedding System for Speech Emotion Recognition

Puneet Kumar[†], Sidharth Jain[‡], Balasubramanian Raman[†], Partha Pratim Roy[†] and Masakazu Iwamura[§]

[†] Department of Computer Science & Engineering

[‡]Department of Civil Engineering

Indian Institute of Technology Roorkee, India

[§]Dept. of Computer Science and Intelligent Systems, Osaka Prefecture University, Japan

{pkumar99@cs, sjain@ce, bala@cs, partha@cs}.iitr.ac.in, masa@cs.osakafu-u.ac.jp

Abstract—In this paper, an end-to-end neural embedding system based on triplet loss and residual learning has been proposed for speech emotion recognition. The proposed system learns the embeddings from the emotional information of the speech utterances. The learned embeddings are used to recognize the emotions portrayed by given speech samples of various lengths. The proposed system implements Residual Neural Network architecture. It is trained using softmax pre-training and triplet loss function. The weights between the fully connected and embedding layers of the trained network are used to calculate the embedding values. The embedding representations of various emotions are mapped onto a hyperplane, and the angles among them are computed using the cosine similarity. These angles are utilized to classify a new speech sample into its appropriate emotion class. The proposed system has demonstrated 91.67% and 64.44% accuracy while recognizing emotions for RAVDESS and IEMOCAP dataset, respectively.

Index Terms—Affective Computing, Deep Learning, Emotion Recognition, End-to-end Speech Processing, Residual Neural Network, Cosine Similarity.

I. Introduction

The need to develop efficient speech processing systems that are capable of recognizing various emotions from the speech is increasing at a fast rate [1]. Such systems are useful for a wide range of applications such as robotics, security, service delivery, language translation, automated identification, intelligent toys, and lie detection [2]. Speech is one of the important ways for a human to portray complex emotions. It can also be used as an efficient method of human-machine interaction [3]. However, a major challenge in human-machine interaction is the correct detection of emotion from speech. It is natural for a human to recognize underlying emotions during their spoken interactions. However, it is difficult for machines to recognize complex emotions in natural speech.

Emotional information included in a speech signal depends on several factors such as speaker, style, language, gender, accent, and sample-duration [3]. The notions of various emotions are highly subjective. People interpret them differently depending upon their culture and environment. Likewise, labeling the speech data with suitable emotion during its preparation is also subjected to human variability. A potential approach to reduce human fluctu-

ations in Speech Emotion Recognition (SER) is to develop an SER system that can recognize speech corresponding to various emotions without human intervention. Such systems are called end-to-end systems, and as opposed to the conventional methods of emotional speech recognition, they do not require manual crafting of acoustic features. There is a need to develop an end-to-end SER system that can learn the emotional patterns in input speech data despite the aforementioned variations and bypassing the intermediate steps of speech processing [4].

The proposed model is based on Residual Neural Network (ResNet) architecture. The model is trained using softmax pre-training and triplet loss function. Then the embedding values are calculated from the weights between the fully connected and embedding layers. The embeddings are mapped onto a hyperplane, and cosine similarity values are calculated by measuring the cosine of the angles among the embedding representations of various emotions. Smaller the angle, higher will be the cosine similarity. The model is trained to minimize the angles among the representations of the speech of similar emotions and maximize the angles among the representations of the speech of different emotions. The computed angles are utilized to classify a new speech sample into its appropriate emotion class. The proposed approach has been validated for two emotional speech datasets - The Ryerson Audio-Visual Database of Emotional Speech and Song (RAVDESS) and The Interactive Emotional Dyadic Motion Capture (IEMOCAP). Recognition accuracies of 91.67% and 64.44% have been observed for RAVDESS and IEMOCAP dataset, respectively.

The major contributions of the paper are as follows. Firstly, a deep neural end-to-end SER system based on triplet loss and residual learning has been proposed. The proposed system is capable of learning emotion-related information from a labeled emotional speech dataset in the form of embeddings. Secondly, the embeddings learned by the proposed system are used to classify the speech samples of various lengths into appropriate emotion classes. Using the embeddings, the proposed system can estimate the emotions in unseen speech utterances.

The rest of the paper is organized as follows. Existing work in the context of speech emotion recognition has been surveyed in Section II. Section III elaborates on the proposed methodology. In Section IV, implementation details and experimental results have been discussed. Finally, Section V concludes the paper and highlights the scope for future research.

II. Related Work

In recent years, several SER approaches have been developed. Feature-based speech recognition systems attempt to extract characteristics from acoustic features such as fundamental frequencies, pitch, prosody, voice quality, Mel frequency cepstrum coefficient (MFCC), and linear prediction cepstrum coefficient (LPCC). They use these features to analyze the emotion patterns of speech samples [2]. In this context, C. Lee et al. [5] used pitch, formants, and speech rate to differentiate positive emotions from negative emotions in speech signals. In another work, J. Rong et al. [6] developed a data pre-processing technique to extract the most relevant acoustic features for emotion recognition. It has been observed that the features derived from high-key emotions such as happiness, anger, and interest show similar properties among themselves which are very different from low-key emotions such as sadness and despair. Hence, there is a need to derive an SER method that is independent of the polarity of the emotional features.

Hidden Markov model (HMM) based statistical methods and support vector machine (SVM) based classifiers have also been used for SER. For example, J. Lorenzo et al. [7] proposed an HMM-based method to detect and alter the emotional context of a speech sample while preserving the identity of the speaker. In another work, P. Shen et al. [8] trained an SVM based classifier to differentiate the emotions present in speech signals based on acoustic features. They used SVM to detect speech emotions accounting for gender-based variations. One of the major problems with using HMM-based SER models is that they are not always able to reliably estimate the parameters of global speech features [2].

The above-mentioned approaches to emotional speech recognition require manual crafting of acoustic features. Hence, it is challenging to come up with an end-to-end SER system using them [9]. Neural network based models are capable of automatically extracting the features from the training data. In context of using them for SER, Stuhlsatz et al. [10] compared the performance of a neural network based Discriminant Analysis with SVM for the classification of emotional speech utterances. The neural network based classifier was observed to outperform the SVM based classifier for speech emotion detection. In another work, Mao et al. [11] used CNN to extract the features from speech spectrograms. Then they classified the features using a binary classifier. Their model outperformed classic machine learning models.

Deep learning based systems have been used for other speech processing tasks as well. For instance, R. Pappagari et al. [12] reused the knowledge learned from speaker recognition for emotion recognition. Similarly, A. Jain et al. [13] implemented speaker independent accent embedding to differentiate multi accent speech. In the context of RNN based SER, N. Majumder et al. [14] implemented attention-based RNN to keep track of the identities of the speakers portraying specific emotions through conversations. In another work, S. Sahoo et al. [15] utilized a pre-trained deep convolutional neural network to predict the emotion classes of the audio segments. Deep neural networks along with residual learning [16] and triplet loss [17] are commonly used for facial expression recognition. They have been used in the area of speech processing also. For example, J. Kim et al. [18] implemented a deep residual network for speech emotion recognition. In another work, H. Bredin [19] used triplet loss along with LSTM to learn the embeddings for speech sequences. The embeddings were later used for speaker identification. As speech emotion recognition is a counterpart of facial emotion recognition, triplet loss and residual learning based techniques may prove to be useful for SER also.

The success of deep neural networks for various speech processing tasks advocates their suitability for SER. However, end-to-end SER using triplet loss and residual learning along with deep neural networks has not been explored to its full potential. With that as an inspiration, various state-of-the-art deep neural architectures have been implemented and the best performing one is implemented in the proposed work. The proposed system also overcomes the issues with existing SER approaches, i.e., need for manual crafting of acoustic features, bias towards the polarity of the emotional features in feature-based SER, and unreliability of statistical SER systems in estimating the parameters of the global speech features.

III. Proposed system

This section describes an end-to-end approach to learn the embedding representations from emotional speech and use them for speech emotion recognition.

A. Problem Formulation

Consider d -dimensional space \mathbb{R}^d where elements in \mathbb{R}^d are represented as $\{x_1, x_2, x_3, \dots, x_d\}$ where x_i is a d -tuple that denotes an embedding vector $f(x) \in \mathbb{R}^d$ mapped from a set of speech utterances $\sum_j y_j$. The projections of such embedding vectors are represented in a hyperplane where emotion similarity is measured using cosine similarity. The objectives of the proposed technique are to:

- i Learn the embeddings from input speech utterances,
- ii Visualize the embeddings projected in a hyperplane to analyze the learned emotion patterns,
- iii Use the learned embeddings to classify an unseen speech utterance into an appropriate emotion class.

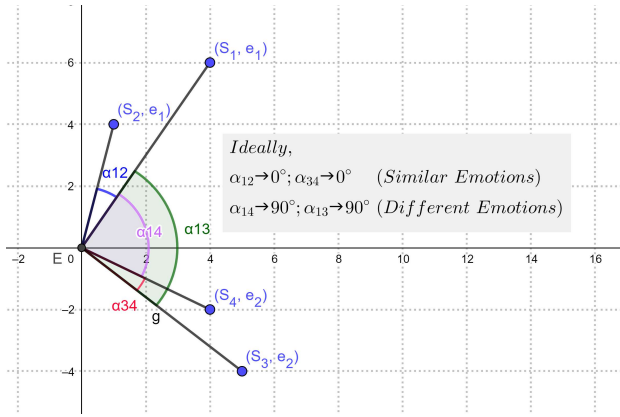


Fig. 1: Hyperplane with projections of speech utterances.




Fig. 2: Representation of the proposed methodology.

Fig. 1 explains the core hypothesis of the proposed work. Ideally, the emotion recognition should be independent of the length of the speech utterance. For example, the speech utterance S_1 is longer than S_2 , but both incorporate the same emotion e_1 . Hence, the angle α_{12} between them is expected to be smaller than the angles between the speech emotions of dissimilar emotions such as α_{13} and α_{14} .

B. Methodology

Various phases of the proposed methodology are represented in Fig. 2 and discussed in the following sections.

1) Phase I: Initialization and Pre-processing: The embeddings are initialized and data is pre-processed.

a. Emotion Embedding Initialization: Emotion embedding is a technique to represent the emotional information of the speech in the form of vectors. It learns the emotional content of the speech and constructs the vector representations for it. The network weights are defined by a temporal average layer which takes frame-level activations and computes the layer activation h as follows:

$$h = \frac{1}{T} \sum_{t=0}^{T-1} x(t) \quad (1)$$

where :

- T : number of frames in a speech utterance.
- t : time instance.
- x : embedding at given time instance t .
- h : activation of the layer.

Emotion embeddings are learned using the weights between the fully connected layer and the embedding layer

after the training completes. Here, initial model weights are initialized using softmax pre-training, and the final weights are assigned to these embeddings in later phases.

b. Audio Cache Generation & Pre-processing: The emotional speech dataset contains audio-clips with emotion labels. For each emotion label, the data is divided into a training set and a testing set. The cache is generated for both, which involves sampling the audio files and trimming the silence. After generating the audio cache, MFCC windows are randomly sampled from it. Since the model has been implemented for constant input shape, an appropriate value for the number of windows for the input signal had to be tuned. We took 10 windows from each input sample using random sampling of 10 continuous MFCC frames. For each window, 39 input values are present which correspond to 13 MFCC values, 13 first-derivative values, and 13 second-derivative values. If the number of windows for an input signal was less than 10, then zero padding was used to keep the input size fixed.

2) Phase II: Embedding Training: A fully connected layer projects the utterance-level representations as embeddings. Emotion characteristics of the speech are learned by training embedding vectors for each emotion. The cross-entropy loss function is used to train the network along with triplet loss.

- Cross-entropy Loss (Softmax Loss): When the output probability of a classification model is between 0 and 1, the cross-entropy loss function can be used to measure its performance. It produces stabler convergence for the data with noisy labels than other methods like mean absolute error loss, and categorical cross-entropy loss [20].
- Triplet Loss: It is an optimization approach that compares a baseline input to a positive input and a negative input. It takes three speech samples and compares them in pairs. The distance between baseline input and positive input is minimized, and the distance between the baseline input and negative input is maximized [17].

$$J = \sum_{i=0}^N F(x_i^{e_1}, x_i^{e_2}, x_i^{e_3})$$

$$F(e_1, e_2, e_3) =$$

$$\max\{d(f(e_1), f(e_2)) - d(f(e_3), f(e_2)) + \alpha, 0\}$$

where :

J : triplet loss cost function.

F : intermediate function.

$x_i^{e_1}$: embeddings for emotion 1.

$x_i^{e_2}$: embeddings for emotion 2.

$x_i^{e_3}$: embeddings for emotion 3.

$d(a, b)$: distance between point a and point b .

f : mapping function for the embeddings.

α : margin between emotion embedding pairs.

(2)

Triplet loss function, J is defined for three emotion embeddings: emotion 1, emotion 2, and emotion 3. Assuming that emotion 1 is similar to emotion 2, but it is dis-similar to emotion 3. Then the distance between e_1 and e_2 will be minimized, and that between e_1 and e_3 will be maximized. This is achieved by function F . Triplets of various emotion pairs are chosen at random, and all the speech utterances are covered. The cost function J computes the overall triplet loss.

Cosine Similarity: Triplet loss is internally optimized using cosine similarity. The cosine similarity has successfully been used to check the similarity among the texts of unequal lengths. With that inspiration, the proposed method has implemented it to check the similarity among speech samples of varying lengths. It calculates the cosine of the angle between the vectors projected in a multi-dimensional space. Smaller the angle becomes, higher will be the cosine similarity. Triplet loss checks the embeddings in pairs. It aims to maximize the cosine similarities for the pairs with the same emotions and minimize those with different emotions

$$\begin{cases} \cos(x_i, x_j) = x_i^T x_j \\ \text{where :} \\ x_i \text{ and } x_j : \text{ two emotion embeddings.} \end{cases} \quad (3)$$

The training process utilizes the above-mentioned concepts, i.e., softmax loss, triplet loss, and cosine similarity. It is carried out in the following two phases.

a. Softmax Pre-training: Softmax pre-training computes both softmax loss and triplet loss, and it trains the model for softmax loss. During pre-training, softmax pre-training is used to initialize the weights of the network. It maps the non-normalized output of the network to a probability distribution over predicted output classes. Softmax pre-training has been found to help avoid getting stuck in a local minimum and producing stabler convergence along with triplet loss [21].

b. Embedding Training with Triplet Loss: This is the full training phase. It computes and trains on both loss values, i.e., softmax loss and triplet loss. Before generating the embeddings, the model weights are L_2 normalized. Then the embedding vector is generated, and triplet loss is calculated and minimized. Although the triplet loss function is well-used for face emotion recognition, it has also found its applications in the area of speech recognition where it has been used to learn and represent the speaker embeddings from speech utterances [21]. Softmax cross-entropy loss works well for a fixed number of classes. However, when there are a variable number of output classes then triplet loss can be used to learn good embeddings for each variation of each class [17].

It is observed that the softmax pre-training along with cross-entropy loss resulted in better performance than triplet loss implementation along with cross-entropy loss.

The combination of triplet loss implementation along with cross-entropy for the input pre-processed with softmax pre-training performed even better.

3) Phase III: Emotion Inference: The steps to infer the emotion category for the unseen samples are as follows.

- a. Compile the test speech utterances in a new folder.
- b. Update the cache for the new folder.
- c. Generate new embeddings using the trained network.
- d. Check cosine similarity for the generated embeddings in comparison with already generated embeddings for the emotional classes in the training phase.
- e. Project various embeddings onto the \mathbb{R}^d hyperplane and measure the angles among them to determine their emotion category.

The model training focuses on learning the emotion embeddings using triplet loss and cosine similarity. The cosine similarity is used for the inference as well. It helped to learn a good representation of the emotional information in speech utterances and gave an idea about the quality of the embeddings produced in the testing phase.

IV. Experiments and results

This section discusses the experimental implementation and analyses the results.

A. Implementation

1) Experimental Setup: Model training has been performed on Nvidia RTX 2070 GPU with 2304 CUDA cores, 288 Tensor cores, and 8 GB Virtual RAM. Model testing has been carried out on Intel(R) Core(TM) i7-7700, 3.70 GHz, 16GB RAM CPU system with Ubuntu 18.04.

2) Dataset Selection and Training Strategy: Speech emotion recognition experiments have been conducted on the following datasets.

- RAVDESS [22] - Ryerson Audio Visual Data of Emotional Speech and Song (RAVDESS) is a standard scripted speech dataset containing eight emotion labels: happy, calm, sad, angry, surprise, fear, neutral, and disgust. Two sentences have been spoken and sung by 24 different speakers. It contains total of 7356 samples.
- IEMOCAP [23] - Interactive Emotional Dyadic Motion Capture (IEMOCAP) is a relatively complex dataset containing impromptu and acted speech samples recorded using emotion sensitive human computer interfaces. It covers 10039 samples labeled with nine different emotion labels. Ten speakers have annotated a total of 12 sentences.

Final implementation has been carried out for complete RAVDESS and IEMOCAP datasets using 70%-30% training-testing split and 10-fold cross-validation. The model training contains two phases: softmax pre-training and triplet loss training. Softmax pre-training has been carried out for 1000 epoch, followed by 1000 epoch of

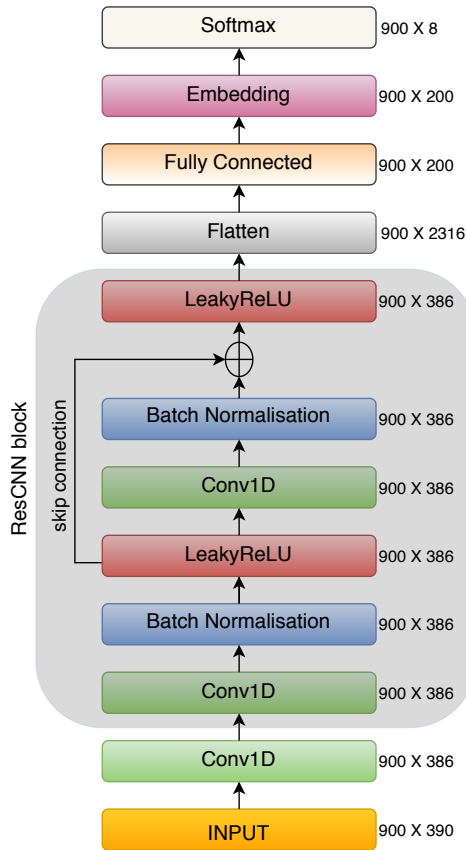


Fig. 3: Schematic description of proposed architecture.

triplet loss fine-tuning. Early stopping has been used to avoid overfitting. The model training is stopped when the accuracy does not improve by at least 0.0005 within 15 epochs. Linear decreasing learning rate ranged from 0.05 to 0.005 has been used. The results have been discussed in Section IV-B.

3) Network Architecture: The ablation study to choose the appropriate network architecture has been performed with a sliced RAVDESS dataset containing 1600 samples, with 200 samples for each emotion class. The dataset was divided into training and testing sets in 70% and 30% ratio. Fully Connected (FC) network, CNN, Residual Neural Network (ResNet), RNN, Long Short Term Memory (LSTM) based RNN, and Gated Recurrent Unit (GRU) based RNN have been evaluated. They have been analyzed to extract emotional features embedded in acoustic input, train utterance level emotion embeddings, and train the network using cosine similarity based triplet loss function. Their details have been presented in Table I. Here ‘x’ represents the total number of layers. The analysis is performed for $x = 6$ to 15.

ResNet with 11 layers (and two residual blocks) performed best among these networks. It has been chosen as the suitable architecture for the proposed implementation. ResNet performs relation extraction using deep residual learning. A residual network contains skip connections

among convolutional layers. It is known to perform better for large networks [21]. Fig. 3 depicts a representative diagram of the implemented ResNet architecture. The network contains the following layers: Input Layer - to feed the input to the network in vector form; Conv1D Layers - to build ResNet block; BatchNorm Layer - to normalize the input and generate triplet loss; LeakyReLU Layer - consists of the activation function to define the output of a particular layer; Flatten Layer - to convert pooled feature map to a single column; Fully Connected Layer - the last layer of the network that takes input from the flatten layer and Softmax Layer - to generate cross-entropy loss. Dimensions of the layers have been shown along with them. The shape of the input is 900X390, where the first dimension 900 corresponds to the batch size, and the second dimension 390 corresponds to the feature values of an individual input speech sample.

B. Results and Discussion

The proposed system learns emotional embeddings from the model weights of the network trained on emotional speech datasets. The learned embeddings are used to recognize the speech samples of various emotions. The cosine similarity values are calculated by applying the Euclidean dot product function onto the embeddings represented in the hyperplane. Triplet loss function is used for model training. It aims to minimize the angles among the representations of the speech of similar emotions and maximize the angles among the representations of the speech of different emotions. The computed angles are utilized to classify a new speech sample into its appropriate emotion class. The following sections discuss the learned embeddings, the angles among their projections, and the emotion classification performance.

1) Emotion Embeddings Visualization: The embedding plots for six important emotions, i.e., anger, neutral, happy, sad, fear, and surprise, have been drawn in Fig. 4. These are the common emotion labels for RAVDESS and IEMOCAP dataset. Hence, they are selected for the embedding plots for comparative understanding. The experimentally obtained embeddings are visualized using t-distributed stochastic neighbor embedding (t-SNE) visualization method.

2) Emotion Classification: The proposed approach has been validated for RAVDESS and IEMOCAP emotional speech dataset. The angles among the embeddings for RAVDESS and IEMOCAP are described in Table II and Table III. The angles for a given speech utterance is checked with all the emotional classes. It is classified into the class with which it makes the least angle.

The proposed approach showed an emotion recognition accuracy of 91.67% for RAVDESS dataset and 64.44% for IEMOCAP dataset. Here, ‘Un-weighted Accuracy’ has been considered, which is defined as total correct predictions over total instances. As discussed in Tables IV

TABLE I: Summary of the ablation study

Architecture	Details	x	Accuracy
FC-x	Fully Connected network with x layers	7	47.22%
CNN-x	Convolutional Neural Network with x layers	8	58.33%
RNN-x	Simple x-layered Recurrent Neural Network	8	55.56%
LSTM-x	x-layered RNN with LSTM units	7	56.94%
GRU-x	x-layered RNN with GRU units	8	54.16%
ResNet-x	Residual Neural Network with x layers	11	61.11%

TABLE II: Angles among various emotional embedding vectors for RAVDESS dataset

	neutral	calm	happy	sad	angry	fearful	disgust	surprise
neutral	0.37°	30.75°	60.34°	89.78°	69.73°	62.63°	80.28°	77.27°
calm		0.51°	78.91°	65.01°	64.61°	66.37°	89.93°	71.71°
happy			0.03°	25.75°	83.22°	66.30°	62.76°	64.39°
sad				0.35°	60.95°	61.52°	63.69°	84.88°
angry					0.63°	84.97°	61.23°	83.40°
fearful						0.03°	67.89°	67.06°
disgust							1.64°	83.63°
surprise								0.50°

TABLE III: Angles among various emotional embedding vectors for IEMOCAP dataset

	anger	sadness	happiness	neutral	excitement	surprise	fear	disgust	frustration
anger	4.51°	51.97°	78.08°	53.39°	56.82°	79.20°	52.68°	62.32°	84.85°
sadness		0.96°	87.23°	70.54°	29.96°	52.28°	67.51°	85.14°	73.09°
happiness			0.99°	67.23°	53.74°	60.81°	77.82°	50.09°	77.06°
neutral				1.12°	56.43°	72.48°	61.06°	45.56°	82.68°
excitement					0.87°	50.74°	68.97°	88.14°	39.29°
surprise						2.66°	83.84°	78.18°	77.65°
fear							0.77°	32.51°	83.56°
disgust								0.73°	73.84°
frustration									2.44°

TABLE IV: Result comparison for RAVDESS dataset

Method	Author	Accuracy
Proposed Method		91.67%
Convolutional Neural Network	M. G. Pinto [24]	91.53%
Artificial Neural Network	K. Tomba et al. [25]	89.16%
Multi Task Hierarchical SVM	B. Zhang et al. [26]	83.15%
Bagged Ensemble of SVMs	A. Bhavan et al. [27]	75.69%
Convolutional Neural Network	D. Issa et al. [28]	71.61%

TABLE V: Result Comparison for IEMOCAP dataset

Method	Author	Accuracy
RNN + Attention	N. Majunder [14]	64.50%
Proposed Method		64.44%
Memory Network	D. Hazarika et al. [29]	63.50%
CNN + Mel Filterbanks	Z. Aldeneh and E. Provost [30]	61.80%
Memory Network	S. Poria et al. [31]	56.13%
CNN + LSTM	J. Zhao [32]	52.14%

and V, the proposed approach has demonstrated comparable performance to the benchmark results. It is to be noted that the performance for RAVDESS and IEMOCAP datasets has been compared for different methods because not all of the available state-of-the-art methods included both datasets in their experiments.

3) Inter and Intra Class Standard Deviation: A given speech utterance is classified into the emotion class with which it makes the least angle. As detailed in Tables VI and VII, deviation in the predicted angles among emotion embedding vectors is observed among similar (intra-class) and different (inter-class) emotion classes. The average intra-class angles are 0.51 and 1.68 for RAVDESS and IEMOCAP. While average inter-class angles are 68.41 and

65.88, respectively. Fig. 5 presents the boxplots for the deviation values. Deviation in intra-class angles affects SER accuracy significantly as it denotes the projections of the embeddings of the speech utterances in the hyperplane. Lesser deviation in the intra-class angles for RAVDESS corresponds to the closeness in their projections and more accurate SER predictions as compared to IEMOCAP.

As observed in Fig. 4, the emotion embeddings for RAVDESS dataset are more clearly defined than IEMOCAP dataset. It is also noted that some overlap has been observed between the embeddings of ‘sad’ and ‘fear’, both of which are negative-valence emotions. On the other hand, both ‘anger’ and ‘surprise’ are emotions with high intensity. They also showed minor overlap between their

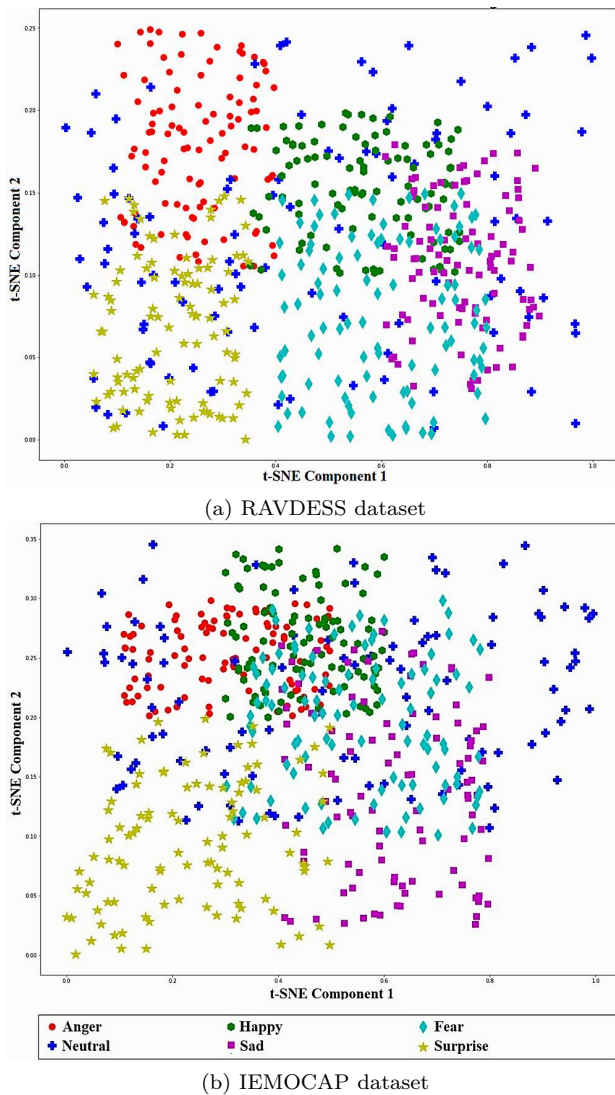


Fig. 4: t-SNE Visualization of Emotion Embeddings.

TABLE VI: Std. Dev. analysis for RAVDESS dataset

	Intra-class deviation	Inter-class deviation
neutral	0.14	2.11
calm	0.00	1.66
happy	0.48	5.32
sad	0.15	4.87
angry	0.12	3.21
fearful	0.47	1.26
disgust	1.13	4.36
surprise	0.01	7.64

embeddings. Another observation is that the network converged faster during the training when softmax pre-training was used along with triplet loss training.

V. Conclusions and future work

In this paper, an end-to-end emotion embedding system has been proposed to learn the emotional patterns from

TABLE VII: Std. Dev. analysis for IEMOCAP dataset

	Intra-class deviation	Inter-class deviation
anger	2.84	0.96
sadness	0.71	1.16
happiness	0.68	3.13
neutral	0.56	1.60
excitement	0.80	9.76
surprise	0.99	3.52
fear	0.91	0.12
disgust	0.94	1.41
frustration	0.77	8.13

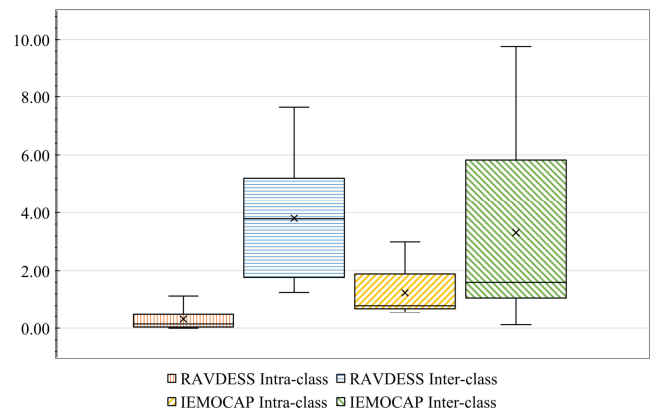


Fig. 5: Boxplots of deviation in predicted angles.

speech samples in the form of an embedding matrix. It projects the mappings of speech information onto a hyperplane where triplet loss is used as a loss function to learn the similarities among various emotions based on cosine similarity. The emotion embedding matrix thus prepared has been used for speech emotion recognition and it demonstrated comparable recognition results to the benchmarks for two different datasets. The proposed system has classified the emotions for RAVDESS dataset with an accuracy of 91.67%, while 64.44% accuracy has been observed for IEMOCAP dataset.

Experiments for various neural architectures to automatically learn the acoustic features were performed, and finally, ResNet was implemented, which demonstrated better performance in finding relational emotion patterns for speech samples. The current implementation requires checking of angles for each speech utterance with each emotion class. This process can be optimized with the aim of reducing training time, computational requirements, and model size. In the future, it is also aimed to use the learned embeddings for other speech processing tasks such as emotional speech synthesis.

Acknowledgements

This research was supported by the Ministry of Human Resource Development (MHRD) INDIA with reference grant number: 1-3146198040.

References

- [1] Z. Zeng, M. Pantic, G. I. Roisman, and T. S. Huang, "A survey of affect recognition methods: Audio, visual, and spontaneous expressions," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 1, pp. 39–58, 2009.
- [2] M. El Ayadi, M. S. Kamel, and F. Karray, "Survey on speech emotion recognition: Features, classification schemes, and databases," *Pattern Recognition*, vol. 44, no. 3, pp. 572–587, 2011.
- [3] T. Vogt, E. André, and J. Wagner, "Automatic recognition of emotions from speech: a review of the literature and recommendations for practical realisation," *Affect and emotion in human-computer interaction*, 1st Ed., pp. 75–91, 2008.
- [4] I. Chiriacescu, "Automatic emotion analysis based on speech," Ph.D. dissertation, M. Sc. THESIS Delft University of Technology, 2009.
- [5] C. M. Lee, S. S. Narayanan et al., "Toward detecting emotions in spoken dialogs," *IEEE Transactions on speech and audio processing*, vol. 13, no. 2, pp. 293–303, 2005.
- [6] J. Rong, G. Li, and Y.-P. P. Chen, "Acoustic feature selection for automatic emotion recognition from speech," *Information processing & management*, vol. 45, no. 3, pp. 315–328, 2009.
- [7] J. Lorenzo-Trueba, R. Barra-Chicote, R. San-Segundo, J. Ferreiros, J. Yamagishi, and J. M. Montero, "Emotion transplantation through adaptation in hmm-based speech synthesis," *Computer Speech & Language*, vol. 34, no. 1, pp. 292–307, 2015.
- [8] U. Jain, K. Nathani, N. Ruban, A. N. J. Raj, Z. Zhuang, and V. G. Mahesh, "Cubic SVM classifier based feature extraction and emotion detection from speech signals," in *IEEE International Conference on Sensor Networks and Signal Processing (SNSP)*, 2018, pp. 386–391.
- [9] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, p. 436, 2015.
- [10] A. Stuhlsatz, C. Meyer, F. Eyben, T. Zielke, G. Meier, and B. Schuller, "Deep neural networks for acoustic emotion recognition: raising the benchmarks," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011, pp. 5688–5691.
- [11] Q. Mao, M. Dong, Z. Huang, and Y. Zhan, "Learning salient features for speech emotion recognition using convolutional neural networks," *IEEE transactions on multimedia*, vol. 16, no. 8, pp. 2203–2213, 2014.
- [12] R. Pappagari, T. Wang, J. Villalba, N. Chen, and N. Dehak, "x-vectors meet emotions: A study on dependencies between emotion and speaker recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7169–7173.
- [13] A. Jain, M. Upreti, and P. Jyothi, "Improved accented speech recognition using accent embeddings and multi-task learning," in *INTERSPEECH*, 2018, pp. 2454–2458.
- [14] N. Majumder, S. Poria, D. Hazarika, R. Mihalcea, A. Gelbukh, and E. Cambria, "Dialoguerrn: An attentive rnn for emotion detection in conversations," in *Proceedings of the Association for the Advancement of Artificial Intelligence (AAAI)*, vol. 33, 2019, pp. 6818–6825.
- [15] S. Sahoo, P. Kumar, B. Raman, and P. P. Roy, "A segment level approach to speech emotion recognition using transfer learning," in *Proceedings of the 5th Asian Conference on Pattern Recognition (ACPR)*, 2019, pp. 435–448.
- [16] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [17] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 815–823.
- [18] J. Kim, G. Englebienne, K. P. Truong, and V. Evers, "Deep temporal models using identity skip-connections for speech emotion recognition," in *Proceedings of the 25th ACM international conference on Multimedia*, 2017, pp. 1006–1013.
- [19] H. Bredin, "TristouNet: triplet loss for speaker turn embedding," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 5430–5434.
- [20] Z. Zhang and M. Sabuncu, "Generalized cross entropy loss for training deep neural networks with noisy labels," in *Advances in Neural Information Processing Systems (NIPS)*, 2018, pp. 8778–8788.
- [21] C. Li, X. Ma, B. Jiang, X. Li, X. Zhang, X. Liu, Y. Cao, A. Kannan, and Z. Zhu, "Deep Speaker: an end-to-end neural speaker embedding system," *arXiv preprint arXiv:1705.02304*, 2017.
- [22] S. R. Livingstone, K. Peck, and F. A. Russo, "RAVDESS: The ryerson audio-visual database of emotional speech and song," in *Annual meeting of the canadian society for brain, behaviour and cognitive science*, 2012, pp. 205–211.
- [23] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "IEMO-CAP: Interactive emotional dyadic motion capture database," *Language resources and evaluation*, vol. 42, no. 4, p. 335, 2008.
- [24] M. G. de Pinto. (2019) Emotion classification RAVDESS. [Online]. Available: <https://github.com/marcogdepinto/Emotion-Classification-Ravdess>
- [25] K. Tomba, J. Dumoulin, E. Mugellini, O. A. Khaled, and S. Hawila, "Stress detection through speech analysis," in *International Joint Conference on e-Business and Telecommunications (ICETE)*, 2018, pp. 560–564.
- [26] B. Zhang, G. Essl, and E. M. Provost, "Recognizing emotion from singing and speaking using shared models," in *IEEE International Conference on Affective Computing and Intelligent Interaction (ACII)*, 2015, pp. 139–145.
- [27] A. Bhavan, P. Chauhan, R. R. Shah et al., "Bagged support vector machines for emotion recognition from speech," *Knowledge-Based Systems*, vol. 184, p. 104886, 2019.
- [28] D. Issa, M. F. Demirci, and A. Yazici, "Speech emotion recognition with deep convolutional neural networks," *Biomedical Signal Processing and Control*, vol. 59, p. 101894, 2020.
- [29] D. Hazarika, S. Poria, R. Mihalcea, E. Cambria, and R. Zimmermann, "Icon: interactive conversational memory network for multimodal emotion detection," in *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018, pp. 2594–2604.
- [30] Z. Aldeneh and E. M. Provost, "Using regional saliency for speech emotion recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 2741–2745.
- [31] D. Hazarika, S. Poria, A. Zadeh, E. Cambria, L.-P. Morency, and R. Zimmermann, "Conversational memory network for emotion recognition in dyadic dialogue videos," in *Conference of the North American Chapter of the Association for Computational Linguistics (ACL)*, 2018, pp. 2122–2132.
- [32] J. Zhao, X. Mao, and L. Chen, "Speech emotion recognition using deep 1D & 2D CNN & LSTM networks," *Biomedical Signal Processing and Control*, vol. 47, pp. 312–323, 2019.