



# Multimodal Emotion Recognition using Transfer Learning from Speaker Recognition and BERT-based models

Sarala Padi<sup>1</sup>, Seyed Omid Sadjadi<sup>1</sup>, Dinesh Manocha<sup>2</sup>, and Ram D. Sriram<sup>1</sup>

<sup>1</sup>NIST, ITL, Gaithersburg, MD

<sup>2</sup>University of Maryland, College Park, MD

## Abstract

Automatic emotion recognition plays a key role in computer-human interaction as it has the potential to enrich the next-generation artificial intelligence with emotional intelligence. It finds applications in customer and/or representative behavior analysis in call centers, gaming, personal assistants, and social robots, to mention a few. Therefore, there has been an increasing demand to develop robust automatic methods to analyze and recognize the various emotions. In this paper, we propose a neural network-based emotion recognition framework that uses a late fusion of transfer-learned and fine-tuned models from speech and text modalities. More specifically, we i) adapt a residual network (ResNet) based model trained on a large-scale speaker recognition task using transfer learning along with a spectrogram augmentation approach to recognize emotions from speech, and ii) use a fine-tuned bidirectional encoder representations from transformers (BERT) based model to represent and recognize emotions from the text. The proposed system then combines the Resnet and BERT-based model scores using a late fusion strategy to further improve the emotion recognition performance. The proposed multimodal solution addresses the data scarcity limitation in emotion recognition using transfer learning, data augmentation, and fine-tuning, thereby improving the generalization performance of the emotion recognition models. We evaluate the effectiveness of our proposed multimodal approach on the interactive emotional dyadic motion capture (IEMOCAP) dataset. Experimental results indicate that both audio and text-based models improve the emotion recognition performance and that the proposed multimodal solution achieves state-of-the-art results on the IEMOCAP benchmark.

## 1. Introduction

As machines and computer-based applications and interactions are continuously progressing, and becoming a part of our daily lives, natural human-computer interaction (HCI) has become increasingly critical [1, 2]. As a result, understanding and responding to the emotional state of individuals is a necessary step in natural HCI [3]. Apart from this, there are other applications where automatic detection and classification of human emotion plays a critical role and can be useful in psychological and nano-physiological studies of human emotional expression. Similarly, emotion information plays an important role in automatic tutoring systems, call center services, gaming, personal assistants, to mention a few. Another potential application can be alerting drivers by detecting their stress or anger levels that could lead to major accidents and impair their driving capabilities [4, 3].

There are a number of modalities with which humans con-

vey their emotions. Examples include speech, text, facial expressions, hand gestures, and so on. Although there are scenarios where facial expressions could be a preferred and more effective way of communicating emotions, speech and text data are abundantly available and more conveniently captured as compared to other modalities. In particular, speech plays a crucial role in conveying the emotional state of a human in the form of prosody and/or paralinguistic context. For speech emotion recognition (SER), traditionally, machine learning (ML) models were developed using engineered features such as mel-frequency cepstral coefficients (MFCC), Chroma-based features, pitch, energy, entropy, and zero-crossing rate [5, 6, 7], to mention a few. However, the performance of such ML models depends on the type and diversity of the features used. Although it remains unclear which features correlate most with various emotions, the research is still ongoing to explore additional features and new algorithms to model the dynamics of feature streams representing human emotions. On the other hand, deep learning-based models can directly learn the task-relevant features from spectrograms or raw waveforms [8, 9, 10, 11], thereby obviating the need for extracting a large set of engineered features [12]. Recent studies have proposed the use of convolutional neural network (CNN) models combined with long short-term memory (LSTM) built on spectrograms and raw waveforms, showing improved SER performance [9, 10, 12, 13, 14, 15, 16]. However, building such complex systems requires large amounts of labeled training data. Also, insufficient labeled training data can potentially make the models overfit to specific data conditions and domains, resulting in poor generalization performance on unseen data conditions.

Another important modality for conveying and capturing emotions is text. To represent and model the textual data for solving natural language processing (NLP) related tasks, different approaches have been developed. The word and sentence embeddings have emerged as the most effective representations and have shown breakthroughs in improving the performance of deep learning models for various NLP applications. The variations of such embeddings include Word to Vector (Word2Vec) [17], Global Vectors (GloVe) [18], and bidirectional encoder representations from transformers (BERT) [19]. Although these methods provide a compact representation of textual data, there are some limitations and challenges in effectively and efficiently leveraging these methods for various tasks. The main challenge is that building such models requires large amounts of data for the task at hand. To overcome this problem, transfer learning and fine-tuning the available pretrained models have shown promise. Nonetheless, the pretrained models are typically trained and tuned for out-of-domain tasks. Accordingly, the tweaked features may not perform well on target tasks such as emotion recognition, where the vocabulary used

to extract the emotion is different from that used for the source task.

This paper presents a multimodal framework for emotion recognition from speech and text using transfer learning and fine-tuning of speaker ResNet and BERT-based models. This paper leverages the advantages of a BERT-based model combined with acoustic feature representations to improve emotion recognition performance. In particular, we use a BERT model to not only extract sentence embeddings from the text but also to fine-tune the model to improve the emotion recognition performance. In addition, to address the data scarcity limitations in speech-based emotion recognition, we use a transfer learning approach combined with a spectrogram augmentation strategy. Specifically, we re-purpose a ResNet [20] model developed for speaker recognition using large amounts of speaker labeled data and use it as a feature descriptor for SER. The model includes a statistics pooling layer that enables the processing of variable-length segments without a need for truncation. Also, we increase the training data size by generating more data samples using spectrogram augmentation[21]. The proposed framework effectively leverages the complementary information from text and speech modalities and combines the scores produced to improve emotion recognition performance. We evaluate the effectiveness of our proposed systems on the interactive emotional dyadic motion capture (IEMOCAP) dataset [22] using speech-only, text-only, and multimodal settings and show that the fusion of the two complementary modalities results in state-of-the-art emotion recognition performance.

## 2. Related Work

In recent years, deep neural network based systems have shown a tremendous success in recognizing the emotions from the speech signal [24, 16, 25]. Specifically, techniques such as bidirectional LSTMs (BLSTM) [23, 15, 24, 25] and time-delay neural networks (TDNN) [26], which can effectively model relatively long contexts compared to their DNN counterparts, have been successfully applied for SER. Nevertheless, as discussed previously, the lack of large amounts of carefully labeled data for building complex models for emotion classification remains a main challenge in SER [27]. To address this, data augmentation methods are used to generate additional training data by perturbing, corrupting, mimicking, and masking the original data samples to enable the development of complex ML models. For example, [16, 28, 29] applied signal-based transformations such as speed perturbation, time-stretch, pitch shift, as well as added noise to original speech waveforms. One disadvantage of these approaches is that they require signal-level modifications, thereby increasing the computational complexity and storage requirements of the subsequent front-end processing. They can also lead to model overfitting due to potentially similar samples in the training set, while random balance can potentially remove useful information [28]. Furthermore, retaining performance on relatively clean conditions has shown to be a challenge while using signal-based transformations [29].

Another effective way to address data scarcity is transfer learning [30, 31, 32]. Transfer learning can leverage the information and knowledge learned from one related task and domain to another, as long as the input remains the same. Several studies have shown that transfer learning approaches outperform prior methods in recognizing emotions even for unseen scenarios, individuals, and conditions [33]. It was also shown that transfer learning increases the feature learning abilities which leads to improved SER performance [34, 35, 36, 37].

However, transfer learning methods have not been fully explored and analyzed for emotion recognition. Particularly, it is unclear whether and how ML models trained for other data-rich speech applications such as speaker recognition would perform for SER. We recently proposed an SER framework based on transfer learning from speaker recognition along with data augmentation [38] which achieved promising improvements in SER performance on the IEMOCAP. In this study, we explore the impact of full network fine-tuning (as opposed to only training the fully connected classification layers) combined with transfer learning on SER performance.

As for text based approaches, word embeddings have emerged as the de facto representation in many natural language processing tasks. The most widely used word embedding models are Word2Vec [17] and GloVe [18]. Both these models are unsupervised and have shown great success for various NLP tasks including sentiment analysis, document indexing, and topic model analysis. However, a major limitation of these models is that the word order is not taken into account which leads to a loss of syntactic and semantic understanding of words in sentences. In addition, fine-tuning these word embeddings for different target tasks and applications is not feasible. To circumvent these limitations, BERT models are used for extracting representations from text data to capture the context. Several studies have successfully combined BERT-based embeddings with the speech-based representations to improve the emotion recognition performance [39, 40, 26]. However, BERT-based embeddings are context dependent and the representations extracted using the pretrained models may not capture the salient domain-dependent information for emotion recognition. This study investigates fine-tuning of a BERT model for text-based emotion recognition

In recent years, ML approaches have been explored where models are trained using multiple modalities and scores are combined by adding an extra layer (i.e., a meta learning layer) either by concatenating (stacking) the embeddings or scores, or using a weighted sum of the embeddings or scores [41, 42, 43]. Multimodal processing and recognition have shown promise in many practical machine learning applications [44, 45]. Since spoken data is composed of audio and text content, prior studies have also investigated the combination of acoustic features/representations and linguistic information for emotion recognition. Particularly, multimodal emotion recognition models have been developed from audio, video, images, or textual data. The models have been combined either at the feature/representation level (aka the early fusion) or at the decision/score level (aka the late fusion) [46, 26, 39]. In this study, we investigate the latter within a multimodal emotion recognition framework using speech and text models.

## 3. Proposed Multimodal Framework

Figure 1 shows the block diagram of the proposed multimodal emotion recognition framework. The following subsections provide a brief description of the speech and text-based models and systems used in this framework.

### 3.1. Text system

#### 3.1.1. BERT

BERT is a language model trained on large amounts of text data that has achieved state-of-art results on many NLP tasks. Because of its bidirectionality and use of transformers with the attention mechanism and positional encoding [47], the BERT

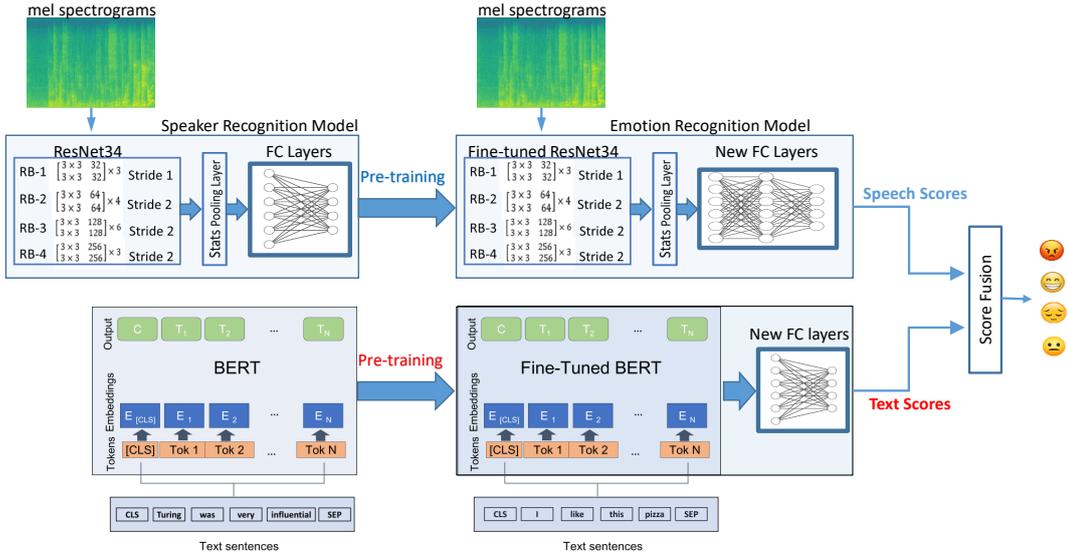


Figure 1: Block diagram of the proposed multimodal emotion recognition framework.

model learns the contextual information from input sentences. The BERT models are context-dependent. So, to use these models for downstream tasks, fine-tuning is required to generate the vectors/embeddings based on the word context. Unlike GloVe, it represents the input as subwords and learns sub-word embeddings (as opposed to whole word embeddings). By using the sub-word representation, the BERT approach overcomes the out of vocabulary limitation of character (ELMO) [48] and word-based representations (GloVe or Word2Vec) [17, 18].

### 3.1.2. Transfer learning

In the BERT setup, preprocessing of the input text is performed to generate tokens and are mapped to indices. As shown in Figure 1, there are two additional tokens called the [CLS] classification token and the [SEP] separate segment token, which are appended at the beginning and end of each sentence, respectively. First, the embedding layer receives a list of input tokens and generates word embeddings for each token by adding positional embeddings to maintain the word order in a sequence. Next, the output of the embedding layer is fed to the multi-head self-attention sublayer to model temporal dependencies. Attention masks are used to exclude the paddings in sequences from attention weight calculations. Finally, we extend the BERT-base model by adding a new classification head, including randomly initialized fully connected (FC) layers with a softmax to get the prediction probabilities for the various emotions. We also fine-tune the network parameters on the training set for the emotion recognition task.

## 3.2. Speech system

Although SER systems traditionally used a large set of low-level time- and frequency-domain features to capture and represent the various emotions in speech, in recent years, many state-of-the-art SER systems have used complex neural network models that learn directly from spectrograms or even raw waveforms. In this study, we build and explore an end-to-end ResNet [20] based system using log-mel spectrograms as input features. ResNet models originally, developed for computer vi-

sion applications [20], have recently gained interest in speech applications such as speaker recognition [49]. We extract high-resolution spectrograms to enable the model to learn the spectral envelope and the coarse harmonic structures for the various emotions.

### 3.2.1. Transfer learning

As noted previously, transfer learning is an ML method where a model initially developed for one task or domain is re-purposed, partly or entirely, for a different but related task/domain. It has recently gained interest in SER [30]. In this study, we re-purpose a model initially developed for speaker recognition to serve as a feature descriptor for SER. We first train a ResNet34 model on large amounts of speaker-labeled audio data, and then we replace the FC layers of the pre-trained model with new randomly initialized FC layers. Finally, we 1) only re-train the new FC layers for an SER task on the IEMOCAP dataset (*linear probing* [50]), and 2) re-train the FC layers and fine-tune the convolutional layers using two different learning rates. As shown in Figure 1, the proposed system employs a statistics pooling layer [51] that aggregates the frame-level information over time and reduces the sequence of frames to a single vector by concatenating the mean and standard deviation computed over frames. The convolutional layers in the ResNet model work at the frame level, while the FC layers work at the segment level.

### 3.2.2. Spectrogram augmentation

There has been recent success in applying a computationally efficient data augmentation strategy, termed spectrogram augmentation, for speech recognition tasks [21]. The spectrogram augmentation technique generates additional training data samples by applying random time-frequency masks to spectrograms to mitigate the over-fitting issue and improve the generalization of speech recognition models. Motivated by promising results seen with the spectrogram augmentation in the speech recognition field, we augment the training data using spectrotemporally modified versions of the original spectrograms. Be-

cause the time-frequency masks are applied directly to spectrograms, we can conveniently apply the augmentation on the fly by eliminating the need to create and store new data files. Similar to the approach taken in [21], we consider two policies, conservative and aggressive, to generate spectrogram augmentations for SER. The parameter settings for the two spectrogram augmentation policies are described in [38].

### 3.3. Multimodal Fusion

To leverage the complementary information captured by the two modalities, we use a weighted average of the scores ( $S_{speech}$  and  $S_{text}$ ) obtained from the speech and text models as follows:

$$S_{fusion} = w_1 \times S_{speech} + w_2 \times S_{text}. \quad (1)$$

The scores are calculated by applying the natural logarithm to outputs of the softmax layer (i.e., class posteriors). Here,  $w_1$  and  $w_2$  are fixed weights assigned to the text and speech modalities, respectively. We constrain the weights to sum to 1, i.e.,  $w_2 = 1 - w_1$ . The weights determine the degree to which each modality contributes to the final score. For emotion classification, we select the emotion category with the highest score.

We investigate two approaches for finding the fusion parameters. In the first approach, we directly use the raw scores from each modality and search for the best setting for  $w_1$  on a hold-out set ( $w_2$  is simply  $1 - w_1$ ). In the second approach, we first pre-process the scores for each modality to have zero mean and unit variance as follows:

$$S_n = \frac{S - \mu_S}{\sigma_S}, \quad (2)$$

where the normalization statistics  $\mu_S$  and  $\sigma_S$  are calculated on a hold-out set. The final fusion score is then obtained by averaging the normalized scores using  $w_1 = w_2 = 0.5$  (i.e., equal weight fusion).

## 4. Experimental Setup

### 4.1. Dataset

We evaluate the effectiveness of the proposed multimodal emotion recognition system on the IEMOCAP dataset [22], which contains improvised and scripted multimodal dyadic conversations between actors of opposite genders. It has 12 hours of speech data from 10 subjects and is pre-segmented into short-cuts. It includes nine categorical emotions and 3-dimensional labels. Our experiments include categorical emotions and speech segments with majority labeling, where at least two annotators agree in the annotation labeling. To replicate the experimental protocols used in many prior studies, we consider four categorical emotions: “angry”, “happy”, “neutral”, and “sad” where we merged “happy” and “excited” into one category and the number of examples per category is 1103, 1636, 1708, and 1084, respectively. Table 1 summarizes the data statistics in the IEMOCAP dataset for the experimental setup considered in this study.

The IEMOCAP dataset comprises five sessions, and the speakers in the sessions are non-overlapping. Therefore, there are ten speakers in the dataset, i.e., five female and five male speakers. To conduct the experiments in a speaker-independent fashion, we use a leave-one-session-out (LOSO) cross-validation strategy, which results in 5 different train-test splits/folds. For each fold, we use the data from 4 sessions for

training and the remaining session for model evaluation. Since the dataset is multi-label and imbalanced, besides the overall accuracy and termed weighted accuracy (WA), we report the average recall over the different emotion categories, termed unweighted accuracy (UA), to present our findings. To understand and visualize the performance of the proposed system within and across the various emotion categories, we compute and report confusion matrices for the experiments. We use similar setups for the speech-only, text-only, and combined model.

Table 1: *Data statistics for the various emotion classes in the IEMOCAP dataset.*

Emotion	Number of Examples
Angry	1103
Happy+Excited	1636
Neutral	1708
Sad	1084
<b>Total</b>	<b>5531</b>

### 4.2. Configuration

#### 4.2.1. Speech system

For speech parameterization, we extract high resolution 128-dimensional log-mel spectrograms from 25 ms frames at a 100 Hz frame rate (i.e., every 10 ms). For feature normalization, we apply a segment level mean and variance normalization<sup>1</sup> which is not ideal because we applied a typical normalization at the recording/conversation level. We have found that normalizing the segments using statistics computed at the conversation level significantly improves the SER performance on the IEMOCAP. However, this violates the independence assumption for the speech segments and did not consider in this study. For the front-end processing, including feature extraction and feature normalization, we use the NIST speaker and language recognition evaluation (SLRE) [52, 53] toolkit. While training the model, we select  $T$ -frame chunks using random offsets over original speech segments, where  $T$  is randomly sampled from the set  $\{150, 200, 250, 300\}$  for each batch. We apply signal padding for speech segments shorter than  $T$  frames. While evaluating the model, we feed the entire duration of the test segments because the statistics pooling layer enables the model to consume variable-length inputs.

As noted previously, the proposed end-to-end SER system uses a pre-trained ResNet34 model built on a speaker recognition task. We train the ResNet34 model with a 512-dimensional embedding layer on millions of speech samples from over 7000 speakers available in the VoxCeleb corpus [54]. To build the speaker recognition model, we apply the same front-end processing described above to extract high-resolution log-mel spectrograms from VoxCeleb data. We conduct experiments using models with and without transfer learning and spectrogram augmentation. For each original speech segment, we generate and augment two spectro-temporally modified versions according to the augmentation policies defined in [38] for both speaker and speech emotion recognition systems during training. We also evaluate our models with and without the statistics pooling layer to study their impact on the emotion recognition task. We use a categorical cross-entropy loss as the objective function to

<sup>1</sup>No voice activity detection (VAD) is applied before feature normalization because it was found to be detrimental to SER performance.

Table 2: Performance comparison of our proposed approach with prior methods that use the LOSO strategy for experiments (Audio, Text, and Audio+Text) on the full IEMOCAP dataset (i.e., both the improvised and scripted portions). Abbreviations: A-Angry, H-Happy, N-Neutral, E-Excited, H+E: Happy and Excited merged. Bold letters indicate the best-performing model. Blanks (–) indicate unreported values.

Modality	Approach	UA [%]	WA [%]
Text	Haiyang Xu et.al [46]	60.3	54.8
	Recurrent encoder [55]	63.5	–
	BERT+self-attention [26]	58.53	59.20
	Glove+self-attention [26]	61.27	62.27
	BERT [39]	55.2	–
	BERT+CNN [40]	66.1	67.0
	<b>Proposed</b>	<b>70.33</b>	<b>70.24</b>
Speech	BLSTM+transfer learnig [34]	51.86	50.47
	VGG19+GAN augmentation [28]	54.6	–
	Recurrent encoder [55]	54.6	–
	CNN+multi-task learning [13]	–	56.10
	BLSTM+attention [15]	58.8	63.5
	CNN+transfer learning [56]	59.54	–
	ResTDNN+attention [26]	61.32	60.64
	LSTM+attention [46]	63.4	57.4
	TDNN+attention [26]	60.64	61.32
	wav2vec-base [57]	63.43	–
	wav2vec-large [57]	65.64	–
	HuBERT-base [57]	64.92	–
	HuBERT-large [57]	67.62	–
	WavLM-base [58]	65.94	–
	WavLM-large [58]	70.03	–
	<b>Proposed (w/o fine-tuning)</b>	<b>64.28</b>	<b>63.74</b>
	<b>Proposed (w/ fine-tuning)</b>	<b>65.97</b>	<b>65.4</b>
Speech + Text	Recurrent encoder [55]	71.8	–
	LSTM+attention+Concat [46]	70.4	69.5
	Audio+GloVe [26]	65.53	66.43
	BERT+CNN Fusion [39]	65.1	–
	Wav2vec+CNN+BERT [40]	73.0	73.5
	<b>Proposed (<math>w_1 = 0.94</math>)</b>	<b>76.07</b>	<b>75.76</b>
<b>Proposed (<math>w_1 = w_2 = 0.5</math>)</b>	<b>75.97</b>	<b>75.01</b>	

train the models. The number of channels in the first block of the ResNet model is set to 32. Pytorch<sup>2</sup> is used for model implementation, the stochastic gradient descent (SGD) optimizer with momentum (0.9), and a batch size of 32. For transfer learning with frozen convolutional weights, an initial learning rate of  $10^{-1}$  is used. To fine-tune the convolutional layers, we use a learning rate of  $10^{-3}$ . The learning rate remains constant for the first 8 epochs and is halved for every other epoch. We use parametric rectified linear unit (PReLU) activation functions in all layers (except for the output) and utilize layer-wise batch normalization to accelerate the training process and improve the generalization properties of the model.

#### 4.2.2. Parameter settings for text model

The most commonly used pre-trained BERT models are BERT-base and BERT-large. The former has 12 encoders with 12 bidirectional self-attention heads, while the latter has 24 encoders with 16 bidirectional self-attention heads. These models were trained on unlabeled data collected from the BooksCorpus (800

<sup>2</sup><https://github.com/pytorch/pytorch>

million words) and English Wikipedia (2500 million words). In our experiments, we use a pretrained BERT-base model, specifically the BERT-base-uncased model with 12-layered transformer blocks with each block containing 12-head self-attention layers and 768 hidden outputs. This model has 110 million parameters in total. We use randomly initialized fully connected layers with  $768 \times 4$  parameters along with softmax, and fine-tune the model for emotion recognition on the IEMOCAP using a categorical cross-entropy loss as the objective function, an Adam optimizer with a learning rate of  $2 \times 10^{-5}$ , and a batch size of 32. We use the Pytorch framework for the BERT model training.

## 5. Performance Evaluations

### 5.1. Results

Table 2 presents the performance comparison of the proposed speech-only, text-only and multimodal systems with several prior approaches for the experimental setup described in Section 4. The fine-tuned BERT model outperforms prior text-based emotion recognition methods on the IEMOCAP with 70.33% UA and 70.24% WA, respectively. The fine-tuned ResNet34 model described in Section 3.2 performs favorably compared to all SER approaches considered in this paper, except for HuBERT-large and WavLM-large with nearly 300 million parameters that use head-to-toe fusion of embeddings from 24 encoder layers. The proposed SER system comfortably outperforms systems that use large number of engineered features [13, 15, 59]. We refer the reader to [38] for the results from the ablation study of the various components (i.e., transfer learning, data augmentation, statistics pooling) of the proposed SER system. In case of multimodal emotion recognition from speech and text, our proposed system outperforms prior methods and achieves state-of-the-art results using both fusion strategies. We note that combining the speech and text-based emotion recognition systems provides remarkable improvements with approximately 6% absolute improvement compared to the best performing unimodal system (76.07% vs 70.33% UA).

### 5.2. Error Analysis

To visualize the performance of the proposed system within and across different emotion categories, confusion matrices for the two modalities as well as their fusion are shown in Figure 2. It is observed from Figure 2(a) and (b) that the system confuses the “happy” class (H+E) with the “neutral” class (N) quite often, while performing the best on the “angry” emotion (A). This is consistent with observations reported in other studies on IEMOCAP [13, 55]. Our informal listening experiments confirm that the “happy” and “neutral” classes are indeed confusable emotion pairs in the IEMOCAP dataset. Combining the Resnet34 and BERT model scores further improves the performance. Moreover, we observe that multimodal fusion reduces the confusion between “Neutral” (N) and “Happy” (H+E) emotions, performing best for “angry” class and achieving accuracies of over 70% for the other three emotion categories.

### 5.3. Cross comparisons

All studies referenced in Table 2 adopt the LOSO strategy to conduct experiments on speech, text, and their fusion. There are other related studies in the literature that either only use the improvised portion of the IEMOCAP dataset or performed experiments using single fold as opposed to 5-fold cross validation

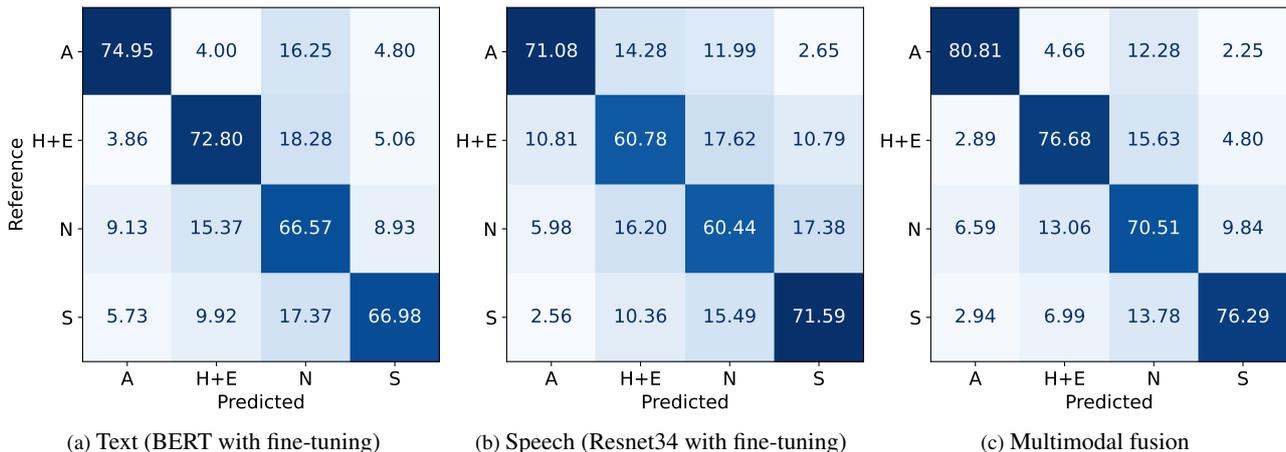


Figure 2: Confusion matrices of the proposed approach for the speech-only, text-only, and multimodal experiments conducted in this study using the LOSO strategy. Abbreviations: A-angry, E-excited, H-happy, N-neutral, S-Sad, and H+E- Happy and Excited merged. Note that, we reported the average of performance metrics (WA & UA) over the five session-wide folds as the result.

Table 3: Cross comparison of our proposed approach with prior methods that use the LOSO strategy for experiments (speech- and text-only) on the full IEMOCAP dataset. ‘M’ indicates million. Bold letters indicate the best-performing model and blanks (–) indicate unreported values.

Method	modality	UA [%]	Parameters
wav2vec-Base[57]	Audio	63.43	95.04M
wav2vec-Large[57]	Audio	65.64	317.38M
HuBERT-Base[57]	Audio	64.92	94.68M
HuBERT-Large [57]	Audio	67.62	316.61M
WavLM-Large[58]	Audio	70.03	316.62M
<b>Proposed</b>	Audio	<b>65.97</b>	<b>21.5M</b>
BERT+attention[26]	Text	58.53	–
BERT+CNN [40]	Text	66.1	–
<b>Proposed</b>	Text	<b>70.33</b>	<b>110M</b>

[60, 9, 16]. On the other hand, in our experiments, we use both the improvised and scripted portions of the IEMOCAP, which is approximately twice the size of the improvised portion alone. Because the experimental setups and the amount of data used for model training and evaluation in those studies are different than ours, we have not included them in Table 2 for comparison. The emotion recognition performance on the improvised portion is known to be better than that on the full dataset (e.g., see [25, 24, 59, 34, 13]).

The latest HuBERT-base results reported in [57] uses a weighted average of embeddings from all transformer layers to achieve 67.62% UA on the IEMOCAP task. However, the first version<sup>3</sup> of the paper reported 62.94% UA only using the embeddings from the last layer. This is more comparable to our proposed SER system with 65.97% UA. Furthermore, from Table 3, we can notice that the HuBERT-large and WavLM-large models each have more than 300 million parameters as compared to our ResNet34 model with nearly 20 million parameters. Although HuBERT-large and WavLM-large models outperform

<sup>3</sup><https://arxiv.org/pdf/2105.01051v1.pdf>

our SER system, these models are computationally expensive and require large amounts of data for self-supervised training.

Similarly for the text modality, our fine-tuned BERT model outperformed prior methods. Note that although [26] reported a performance of 71.22% UA, the experimental setup uses a context window [-3,3] including 3 preceding and succeeding sentences to generate the emotion prediction score for each sentence. This violates the independence assumption for the segments, hence it is not considered for comparison in this study. In addition, although this setup may be applicable to the IEMOCAP where context is inferred from the data, it may not generalize well to other datasets.

## 6. Conclusions

In this study, we presented a multimodal emotion recognition framework using transfer learning and fine-tuning of pre-trained speaker recognition and BERT models. Specifically, we repurposed a pre-trained ResNet model from speaker recognition that was trained using large amounts of speaker-labeled data. Further, we fine-tuned the pretrained ResNet34 model to extract features from high-resolution log-mel spectrograms to improve the speech emotion recognition performance. In addition, we adopted a spectrogram augmentation technique to generate additional training data samples by applying random time-frequency masks to log-mel spectrograms to mitigate overfitting and improve the generalization of emotion recognition models. We also explored a BERT-based emotion recognition system and fine-tuned the pretrained weights to improve the emotion recognition performance. We further improved the emotion recognition performance by fusing the complementary information available from speech and text modalities. We evaluated the proposed system using speech-only, text-only and multimodal settings and compared the performance of our system against that of several prior state-of-the-art studies. The proposed system consistently provided competitive performance across the two modalities as well as their fusion, achieving state-of-the-art results using text and multimodal settings. Results from this study show that models trained for data-rich applications

such as speaker recognition and language modeling can be repurposed using transfer learning to improve the emotion recognition performance under data scarcity constraints. In the future, we plan to extend our work to other datasets and other languages. We also plan to explore more efficient alternatives to network fine-tuning such as head-to-toe probing [50].

## 7. Disclaimer

The views and conclusions presented in this paper are those of the authors and should not be interpreted as the official findings, either expressed or implied, of NIST or the U.S. Government.

## 8. References

- [1] B. W. Schuller, "Speech emotion recognition: Two decades in a nutshell, benchmarks, and ongoing trends," *Communications of the ACM*, vol. 61, no. 5, pp. 90–99, 2018.
- [2] H. Richard, R. Tom, R. Yvonne, and S. Abigail, *Being Human: Human-Computer Interaction in The Year 2020*. Report, Microsoft Corporation, 2008.
- [3] R. Beale and C. Peter, "The role of affect and emotion in HCI," in *Affect and emotion in human-computer interaction*. Springer, 2008, pp. 1–11.
- [4] R. Cowie, E. Douglas-Cowie, N. Tsapatsoulis, G. Votsis, S. Kollias, W. Fellenz, and J. G. Taylor, "Emotion recognition in human-computer interaction," *IEEE Signal Processing Magazine*, vol. 18, no. 1, pp. 32–80, 2001.
- [5] D. Ververidis and C. Kotropoulos, "Emotional speech recognition: Resources, features, and methods," *Speech Communication*, vol. 48, no. 9, pp. 1162–1181, 2006.
- [6] K. Han, D. Yu, and I. Tashev, "Speech emotion recognition using deep neural network and extreme learning machine," in *Proc. INTERSPEECH*, 2014, pp. 223–227.
- [7] O.-W. Kwon, K. Chan, J. Hao, and T.-W. Lee, "Emotion recognition by speech signals," in *Proc. INTERSPEECH*, 2003, pp. 125–128.
- [8] Z. Yang and J. Hirschberg, "Predicting arousal and valence from waveforms and spectrograms using deep neural networks," in *Proc. INTERSPEECH*, 2018, pp. 3092–3096.
- [9] A. Satt, S. Rozenberg, and R. Hoory, "Efficient emotion recognition from speech using deep learning on spectrograms," in *Proc. INTERSPEECH*, 2017, pp. 1089–1093.
- [10] X. Ma, Z. Wu, J. Jia, M. Xu, H. Meng, and L. Cai, "Emotion recognition from variable-length speech segments using deep learning on spectrograms," in *Proc. INTERSPEECH*, 2018, pp. 3683–3687.
- [11] M. Gao, J. Dong, D. Zhou, Q. Zhang, and D. Yang, "End-to-end speech emotion recognition based on one-dimensional convolutional neural network," in *Proc. ACM ICIAI*, 2019, pp. 78–82.
- [12] P. Yenigalla, A. Kumar, S. Tripathi, C. Singh, S. Kar, and J. Vepa, "Speech emotion recognition using spectrogram & phoneme embedding," in *Proc. INTERSPEECH*, 2018, pp. 3688–3692.
- [13] M. Neumann and N. T. Vu, "Attentive convolutional neural network based speech emotion recognition: A study on the impact of input features, signal length, and acted speech," in *Proc. INTERSPEECH*, 2017, pp. 1263–1267.
- [14] G. Keren and B. Schuller, "Convolutional RNN: an enhanced model for extracting features from sequential data," in *Proc. IEEE IJCNN*, 2016, pp. 3412–3419.
- [15] S. Mirsamadi, E. Barsoum, and C. Zhang, "Automatic speech emotion recognition using recurrent neural networks with local attention," in *Proc. IEEE ICASSP*, 2017, pp. 2227–2231.
- [16] M. Sarma, P. Ghahremani, D. Povey, N. K. Goel, K. K. Sarma, and N. Dehak, "Emotion identification from raw speech signals using DNNs," in *Proc. INTERSPEECH*, 2018, pp. 3097–3101.
- [17] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arxiv.preprint arxiv:1301.3781*, 2013.
- [18] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1532–1543.
- [19] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [20] 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*, 2016, pp. 770–778.
- [21] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," in *Proc. INTERSPEECH*, 2019, pp. 2613–2617.
- [22] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "IEMOCAP: Interactive emotional dyadic motion capture database," *Language resources and evaluation*, vol. 42, no. 4, p. 335, 2008.
- [23] C.-W. Huang and S. Narayanan, "Attention assisted discovery of sub-utterance structure in speech emotion recognition," in *Proc. INTERSPEECH*, 2016, pp. 1387–1391.
- [24] G. Ramet, P. N. Garner, M. Baeriswyl, and A. Lazaridis, "Context-aware attention mechanism for speech emotion recognition," in *Proc. IEEE SLT Workshop*, 2018, pp. 126–131.
- [25] S. Tripathi, T. Sarthak, and H. Beigi, "Multi-modal emotion recognition on IEMOCAP dataset using deep learning," *arXiv preprint arXiv:1804.05788*, 2018.
- [26] W. Wu, C. Zhang, and P. C. Woodland, "Emotion recognition by fusing time synchronous and time asynchronous representations," in *Proc. IEEE ICASSP*, 2021, pp. 6269–6273.
- [27] S. Albanie, A. Nagrani, A. Vedaldi, and A. Zisserman, "Emotion recognition in speech using cross-modal transfer in the wild," in *Proc. ACM ICM*, 2018, pp. 292–301.
- [28] A. Chatziagapi, G. Paraskevopoulos, D. Sgouropoulos, G. Pantazopoulos, M. Nikandrou, T. Giannakopoulos, A. Katsamanis, A. Potamianos, and S. Narayanan, "Data augmentation using GANs for speech emotion recognition," in *Proc. INTERSPEECH*, 2019, pp. 171–175.

- [29] 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 *Proc. IEEE ICASSP*, 2020, pp. 7169–7173.
- [30] K. Feng and T. Chaspari, "A review of generalizable transfer learning in automatic emotion recognition," *Frontiers in Computer Science*, vol. 2, no. 9, 2020.
- [31] J. Gideon, S. Khorram, Z. Aldeneh, D. Dimitriadis, and E. M. Provost, "Progressive neural networks for transfer learning in emotion recognition," in *Proc. INTERSPEECH*, 2017.
- [32] G. Boateng and T. Kowatsch, "Speech emotion recognition among elderly individuals using multimodal fusion and transfer learning," in *Proc. ACM ICMI*, 2020, pp. 12–16.
- [33] J. Gideon, M. McInnis, and E. M. Provost, "Improving cross-corpus speech emotion recognition with adversarial discriminative domain generalization (ADDoG)," *IEEE Trans. Affective Computing*, 2019.
- [34] S. Ghosh, E. Laksana, L.-P. Morency, and S. Scherer, "Representation learning for speech emotion recognition," in *Proc. INTERSPEECH*, 2016, pp. 3603–3607.
- [35] S. Latif, R. Rana, S. Younis, J. Qadir, and J. Epps, "Transfer learning for improving speech emotion classification accuracy," in *Proc. INTERSPEECH*, 2018, pp. 257–261.
- [36] J. Deng, Z. Zhang, E. Marchi, and B. Schuller, "Sparse autoencoder-based feature transfer learning for speech emotion recognition," in *Proc. Humaine Association Conference on Affective Computing and Intelligent Interaction*, 2013, pp. 511–516.
- [37] P. Song, Y. Jin, L. Zhao, and M. Xin, "Speech emotion recognition using transfer learning," *IEICE Trans. Information and Systems*, vol. 97, no. 9, pp. 2530–2532, 2014.
- [38] S. Padi, S. O. Sadjadi, R. D. Sriram, and D. Manocha, "Improved speech emotion recognition using transfer learning and spectrogram augmentation," in *Proc. ACM ICMI*, 2021, pp. 645–652.
- [39] L. Pepino, P. Riera, L. Ferrer, and A. Gravano, "Fusion approaches for emotion recognition from speech using acoustic and text-based features," in *Proc. IEEE ICASSP*. IEEE, 2020, pp. 6484–6488.
- [40] M. R. Makiuchi, K. Uto, and K. Shinoda, "Multimodal emotion recognition with high-level speech and text features," *arxiv,preprint arxiv:2111.10202*, 2021.
- [41] N. Majumder, S. Poria, D. Hazarika, R. Mihalcea, A. Gelbukh, and E. Cambria, "Dialoguerrn: An attentive rnn for emotion detection in conversations," *arxiv,preprint arxiv:1811.00405*, 2019.
- [42] P. Tzirakis, G. Trigeorgis, M. A. Nicolaou, B. W. Schuller, and S. Zafeiriou, "End-to-end multimodal emotion recognition using deep neural networks," *IEEE Journal of Selected Topics in Signal Processing*, p. 1301–1309, Dec 2017.
- [43] H. Ranganathan, S. Chakraborty, and S. Panchanathan, "Multimodal emotion recognition using deep learning architectures," in *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2016, pp. 1–9.
- [44] N. Sebe, I. Cohen, and T. S. Huang, "Multimodal emotion recognition," in *Handbook of pattern recognition and computer vision*. World Scientific, 2005, pp. 387–409.
- [45] S. Haq and P. J. Jackson, "Multimodal emotion recognition," in *Machine audition: principles, algorithms and systems*. IGI Global, 2011, pp. 398–423.
- [46] H. Xu, H. Zhang, K. Han, Y. Wang, Y. Peng, and X. Li, "Learning alignment for multimodal emotion recognition from speech," *CoRR*, vol. abs/1909.05645, 2019.
- [47] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *arxiv, preprint arxiv:1706.03762*, 2017.
- [48] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, "Deep contextualized word representations," *arxiv,preprint arxiv:1802.05365*, 2018.
- [49] H. Zeinali, S. Wang, A. Silnova, P. Matějka, and O. Plchot, "BUT system description to VoxCeleb speaker recognition challenge 2019," *arXiv preprint arXiv:1910.12592*, 2019.
- [50] U. Evci, V. Dumoulin, H. Larochelle, and M. C. Mozer, "Head2Toe: Utilizing intermediate representations for better transfer learning," 2022.
- [51] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *Proc. IEEE ICASSP*, 2018, pp. 5329–5333.
- [52] S. O. Sadjadi, C. Greenberg, E. Singer, D. Reynolds, L. Mason, and J. Hernandez-Cordero, "The 2019 NIST audio-visual speaker recognition evaluation," in *Proc. Speaker Odyssey Workshop*, 2020, pp. 259–265.
- [53] S. O. Sadjadi, T. Kheyrkhan, A. Tong, C. Greenberg, E. Singer, D. Reynolds, L. Mason, and J. Hernandez-Cordero, "The 2017 NIST language recognition evaluation," in *Proc. Speaker Odyssey Workshop*, 2018, pp. 82–89.
- [54] A. Nagrani, J. S. Chung, W. Xie, and A. Zisserman, "Voxceleb: Large-scale speaker verification in the wild," *Computer Speech & Language*, vol. 60, pp. 1–15, 2020.
- [55] S. Yoon, S. Byun, and K. Jung, "Multimodal speech emotion recognition using audio and text," in *Proc. IEEE SLT Workshop*. IEEE, 2018, pp. 112–118.
- [56] M. Neumann and N. T. Vu, "Improving speech emotion recognition with unsupervised representation learning on unlabeled speech," in *Proc. IEEE ICASSP*, 2019, pp. 7390–7394.
- [57] S.-w. Yang, P.-H. Chi, Y.-S. Chuang, C.-I. J. Lai, K. Lakhotia, Y. Y. Lin, A. T. Liu, J. Shi, X. Chang, G.-T. Lin *et al.*, "SUPERB: Speech processing universal performance benchmark," in *Proc. INTERSPEECH*, 2021.
- [58] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *arXiv preprint arXiv:2110.13900*, 2021.
- [59] L. Tarantino, P. N. Garner, and A. Lazaridis, "Self-attention for speech emotion recognition," in *Proc. INTERSPEECH*, 2019, pp. 2578–2582.
- [60] C. Etienne, G. Fidanza, A. Petrovskii, L. Devillers, and B. Schmauch, "CNN+LSTM architecture for speech emotion recognition with data augmentation," *arXiv preprint arXiv:1802.05630*, 2018.