



# SELM: Enhancing Speech Emotion Recognition for Out-of-Domain Scenarios

Hazim Bukhari, Soham Deshmukh, Hira Dharmyal, Bhiksha Raj, Rita Singh

Carnegie Mellon University

{hbukhari, sdeshmukh, hyd, bhikshar, rsingh}@andrew.cmu.edu

## Abstract

Speech Emotion Recognition (SER) has been traditionally formulated as a classification task. However, emotions are generally a spectrum whose distribution varies from situation to situation leading to poor Out-of-Domain (OOD) performance. We take inspiration from statistical formulation of Automatic Speech Recognition (ASR) and formulate the SER task as generating the most likely sequence of text tokens to infer emotion. The formulation breaks SER into predicting acoustic model features weighted by language model prediction. As an instance of this approach, we present SELM, an audio-conditioned language model for SER that predicts different emotion views. We train SELM on curated speech emotion corpus and test it on three OOD datasets (RAVDESS, CREMAD, IEMOCAP) not used in training. SELM achieves significant improvements over the state-of-the-art baselines, with 17% and 7% relative accuracy gains for RAVDESS and CREMA-D, respectively. Moreover, SELM can further boost its performance by Few-Shot Learning using a few annotated examples. The results highlight the effectiveness of our SER formulation, especially to improve performance in OOD scenarios.

**Index Terms:** speech emotion recognition, out-of-distribution, audio language models, few-shot learning

## 1. Introduction

Understanding emotions from spoken language helps humans in decision-making and social interaction. Speech Emotion Recognition (SER) aims to enable machines to detect emotions from human speech. The emotions are commonly represented as either categorical emotions (*Happy, Angry, Sad*), as Sentiments (*Positive, Neutral, Negative*), or as continuous values of valence, arousal, and dominance ratings. Automatically detecting emotions from speech has multiple applications in various domains such as interactive voice assistants, caller-agent conversation analysis, health care, education, and entertainment.

Literature on SER formulates it as a classification task [1]. The model architecture consists of pretrained speech encoder like Wav2vec2 [2], HuBERT [3], and Whisper [4] followed by a classifier. The pretrained speech encoder is usually completely or partially frozen and the classifier is trained from scratch on the target data. This model architecture performs well on in-domain data [5]. However, the performance drops drastically on out-of-domain data. Moreover, the learned classifier is static and only predicts predefined emotions and hence needs to be retrained for every domain and target task. The recent works in audio-text learning [6, 7, 8, 9, 10] aim to avoid training classifiers for each target task and domain. To remove the fixed classifier stage, audio-text models use cosine similarity between audio embeddings and classes to determine predictions where the classes are represented by text embeddings. The popular instances of such

approaches like CLAP [7], Pengi [6], LTU [8, 9] allow users to define emotion classes to predict at test-time. Moreover, as emotion is a spectrum, flexible class prediction enables users to define emotion as “happy but little sad” and “initially happy tone but overall sad” or “happy initially and sad later”. Though audio-text approaches are trained on large-scale audio data, they still cannot generalize to OOD scenarios, where the speech data differs from the training data in terms of speakers, languages, accents, recording conditions, or emotional categories. For example, *Happy* emotion in call-center has a significantly different energy than *Happy* in business meetings, leading to misclassifications. This reduces the applicability of audio-text models for SER in-the-wild.

Few-Shot Learning (FSL) approaches utilize a few annotated examples from the target domain to adapt the pre-trained model. This allows SER models to be adapted to different domains like call-center, and health care with few examples and minimal updates. Two popular FSL approaches for SER are LanSER [11] and AudioFlamingo [12]. LanSER utilizes a pre-trained Large Language Model (LLM) through weakly-supervised learning and then uses 10% of data to improve OOD performance. Similarly, AudioFlamingo takes the approach of building a strong foundation audio model that can be later adapted to OOD dataset using as less as 4-shots. However, there exists a gap in in-domain performance and the OOD performance of AudioFlamingo and LanSER. Moreover, LanSER requires large amount of annotated audio data to improve OOD performance. For example, 10% for CREMAD [13] is 744 audio files, which is costly to annotate.

In this work, we provide an alternate formulation for SER and propose SELM, a speech emotion recognition model that improves OOD performance and is adaptable to new test distribution with minimal or few-annotated examples. SELM uses an audio-conditioned language modeling approach for SER where the prediction of different emotion views is formulated as a text generation task. Our contributions are as follows: (1) An alternate formulation of the task of SER which breaks SER into Acoustic Model and Language Model. (2) Introduce SELM which is an instance of the above formulation and provides SoTA performance in Out-of-Domain scenarios. (3) Propose Few-Shot Learning approach for SELM (4) Extensively testing SELM in various setups: In-domain, OOD, and Few-Shot Learning on three public datasets (RAVDESS, CREMA-D and IEMOCAP), thus establishing a baseline for future work.

## 2. Formulation

An ASR system produces the most likely word sequence for a speech signal. The statistical approach to ASR has been:

$$w^* = \arg \max_w p(x|w)p(w)$$

where  $x$  represents the audio,  $w$  represents a word sequence and the optimal word decoding is  $w^*$ . The  $p(w)$  is modeled by the language model and  $p(x|w)$  is calculated by the acoustic model. Traditionally speech emotion recognition has been formulated as a classification problem. However, one can formulate speech emotion recognition as generating emotion tokens ( $e$ ), conditioned on audio, where the emotion tokens are actually text tokens describing emotion. Taking inspiration from the ASR formulation, we provide the statistical formulation for speech emotion recognition as :

$$e^* = \arg \max_e p(e|x, w)$$

where we want to predict the best emotion sequence given acoustic sequence  $x$  and word sequence  $w$  (transcript). The emotion is considered a sequence of tokens instead of classes. In practice, we use Byte pair encoding used by GPT2 to convert text into tokens. Following Bayes' rule, the posterior probability in the above equation can be expressed as:

$$e^* = \arg \max_e \frac{p(x|e, w)p(e, w)}{p(x, w)}$$

$$e^* = \arg \max_e \frac{p(x|e, w)p(e|w)p(w)}{p(x|w)p(w)}$$

$$e^* = \arg \max_e \frac{p(x|e, w)p(e|w)}{p(x|w)}$$

$p(x|w)$  is constant with respect to ranking of hypothesis for  $e$ , hence can be discarded

$$e^* = \arg \max_e p(x|e, w)p(e|w)$$

The  $p(e|w)$  tells us given the word sequence what is the likely emotion. This can be modeled by a pre-trained Language Model, for example GPT2. The  $p(x|e, w)$  tells us what acoustic features can be observed given the particular emotion and word sequence (transcript). It should contain a representation of the distinct sounds that make up each emotion and word in language model.

### 3. Methodology

We realize the formulation (Section 2) as audio-conditioned language modeling for Speech Emotion Recognition. In this section, we describe SELM (Figure 1), which takes an audio recording and prompt as input and produces text as output.

#### 3.1. Architecture

**Audio encoder.** The audio encoder extracts dense representations from audio. We use Wav2Vec2 as the choice of audio encoder due to its near SOTA performance for SER [5]. The 4th layer of Wav2Vec2 contains the information relevant to the task of SER [5]. Therefore, we extract the 4th layer of Wav2Vec2 as our frozen audio feature extractor.

**Text Embedder.** The text embedder consists of a Language Model tokenizer followed by embedding lookup to convert language model tokens to continuous embeddings. We keep the tokenizer and embedding lookup same as the Language Model used, which in our case is GPT2 [14]

**Audio Projection.** The audio projection consists of two learnable linear layers with GeLU activation in between. This is similar to CLAP projection [7], and performs a non-linear transformation of the 4th layer of Wav2Vec2 hidden state. This transformed representation is passed to the Audio Mapper.

**Audio Mapper.** The audio mapper consists of a sequence mapper and a transformer layer. The sequence mapper is a linear

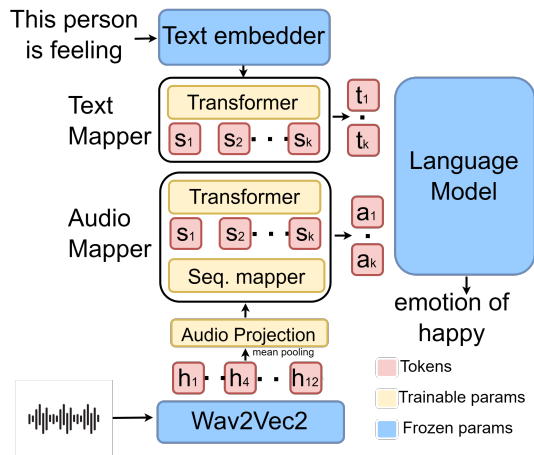


Figure 1: *SELM: Speech Emotion Language Model*. The model is fed with audio and input prompts, which get independently encoded by audio projection & audio mapper, and text embedder. The encoded audio and text is used to prompt a Language Model. In the figure, the input text prompt is “this person is feeling” and SELM outputs “emotion of happy”. The audio projection and audio mapper are learned during training while the Language Model and Wav2vec2 are frozen.

layer followed by a reshape operation to convert a single embedding into a fixed set of continuous embeddings ( $t_1, t_2, \dots, t_k$ ). The sequence is passed to a self-attention Transformer. The transformer learns to map the generated sequence from audio ( $s_1, \dots, s_k$ ) to the latent space of the Language Model ( $a_1, \dots, a_k$ ).

**Text Mapper.** The text mapper consists of a learnable Transformer layer to transform text embeddings into the latent space of the Language Model ( $t_1, t_2, \dots, t_k$ ).

**Language Model.** The latent prefix for the Language Model is the concatenation of the audio prefix ( $a_1, a_2, \dots, a_k$ ) and text prefix ( $t_1, t_2, \dots, t_k$ ). The Language Model generates a text (*emotion of happy*) conditioned on this latent prefix. We use GPT2 [14] as the choice of Language Model and a prefix size of 10 for audio and text respectively.

#### 3.2. Training

The model is trained using the next-token prediction objective. Let one example of training data consist of audio  $x_{a_i}$ , input text prompt  $x_{t_i}$  and ground truth response  $y_{t_i}$ . The audio encoder and audio mapper transform the audio  $x_{a_i}$  into a continuous sequence of embeddings ( $a_{1_i}, a_{2_i}, \dots, a_{k_i}$ ), where  $k$  is the prefix size. Similarly, the text embedder and text mapper transforms the prompt  $x_{t_i}$  into a continuous sequence of embeddings  $t_{1_i}, t_{2_i}, \dots, t_{k_i}$ . Both sequences are concatenated to form the prefix  $z_i$ :

$$z_i = a_{1_i}, a_{2_i}, \dots, a_{k_i}, t_{1_i}, t_{2_i}, \dots, t_{k_i} \quad (1)$$

Then the Language Model conditioned on  $z_i$  produces a sequence of output tokens  $o_{t_i}$ . The model is trained to predict the next text tokens  $o_{t_i}$  ( $t \in [0, l]$ ) conditioned on  $z_i$  in an autoregressive fashion. The loss function is Cross-Entropy:

$$\mathcal{L} = - \sum_{t=1}^l \log p_{\theta}(y_{t_i} | z_i, y_{0_i}, \dots, y_{l_i}) \quad (2)$$

where  $\theta$  denotes the model’s trainable parameters which consist of Audio Projection, Audio Mapper, and Text Mapper. We keep the Wav2Vec2, Text embedder, and Language Model frozen.

### 3.3. Inference

The test audio and test input prompt are used to build prefix  $z_i$  to prompt the Language Model. The Language Model predicts the next token based on beam search decoding with a beam size of 3. The output of the model is a text and can be directly used in any downstream pipeline

For evaluating metrics, the free-form text answer from SELM has to be converted to 1 out of C classes for each dataset or user-specified classes. For example, SELM generates the emotion as frustration but the user wants classification into 4 classes of *Happy, Sad, Angry, Neutral*. Therefore, the generated text needs to be mapped into one of the C classes. For this, we encode the generated text and classes with CLAPS text encoder. Then use cosine similarity between text embedding of generated text and text embedding of classes to determine the most likely class.

## 4. Experimental setup

### 4.1. Datasets

**Training Datasets:** The training dataset for SELM consists of triplets of form (audio, prompt, output). Training on all emotion views is proven to improve performance [15]. Hence, we train SELM to predict all three views of emotions: categorical, sentiment, and dimensional. For categorical data, we source audio and labels from MSP Podcast [16], CMU MOSI [17], CMU MOSEI [18], MELD [19] and use "This person is" as the prompt. For sentiment data, we convert categorical labels from MSP Podcast and MELD to sentiment with "This sentiment is" as a prompt. While, for dimensional scores, we utilize the labels available from the MSP Podcast and prompt with "Describe emotion parameters". In total, we train on 315k triplets.

Dataset	Files	Dur. (secs)	Output	Setup
RAVDESS	2.5k	$\leq 5$	MC (8)	5 folds
CREMA-D	7k	5	MC (6)	5 folds
IEMOCAP	2	4.5	MC (6)	train/test

Table 1: *The datasets used to simulate OOD setup. They differ in terms of speakers, accents, recording conditions, and emotional categories. For the "Output Type" column, MC to multiclass and (.) represents the classes.*

**Evaluation Datasets:** We use three datasets for evaluating SER performance: RAVDESS [20], CREMAD [13], IEMOCAP [21]. These datasets differ in the number of emotional classes, the type of features (audio-only vs. multimodal), and their specific data collection setup (Table 1). The above datasets are not used in training SELM. This makes the three datasets ideal for simulating OOD setup and for benchmarking.

### 4.2. Setup

**In-Domain.** To check the in-distribution performance of SELM we use the in-domain setup. In this setup, the model is trained using the training subset of the Evaluation Datasets and evaluated on the testing subset of them. In explicit train-test subset are not available, we perform N-fold cross-validation.

**Out-of-Domain.** We simulate the Out-of-Distribution setup by excluding some datasets from training. Therefore, the model is evaluated on three Evaluation Datasets namely RAVDESS, CREMA-D, and IEMOCAP, not used in training.

**Few-Shot Learning.** Adapting to new distribution requires annotated examples for the target domain. To test the efficiency of this adaptation, we explore Few-Shot Learning (FSL) for SELM.

For the FSL setup, we inherit the model pretrained SELM model. Then the model is finetuned using N examples from target domain. The finetuning is restricted to specific parts of model to be parameter efficient and to prevent loss of generalization ability of SELM. We perform FSL in two settings, 4-shot and 8-shot, where 4 examples are 8 examples per class are randomly picked from the training set to finetune specific parameters of SELM.

## 5. Results

Unweighted accuracy is used as the primary evaluation metric. The experiments in this section are designed to test SELM in different scenarios, namely: In-domain (Section 5.1), Out-of-Domain (Section 5.2), Few-Shot (Section 5.3). The Section 5.4 contains ablation studies to determine which parameters to update for Few-Shot Learning.

### 5.1. In-Domain

In this Section, we compare our formulation of SER with the traditional classification-based SER. In SELM, the audio encoder is Wav2Vec and is frozen i.e. not learned during training. We compare against literature baselines where the Wav2Vec2 is frozen followed by learning single or multiple linear layers. This experiment is called In-Domain setup as the model is trained on the train subset of the dataset followed by testing on the test subset. The results are shown in Table 2 for three datasets. We present SELM performance in the first row, and compare it to the methodology where Wav2Vec2 features are used followed by a simple neural network on top. We believe this is a fair comparison as the acoustic model architecture is the same in both models. SELM's better performance shows the benefit of decomposing the SER task into an Acoustic Model followed by a Language Model reweighting.

Model	In-domain Dataset		
	RAVDESS $\uparrow$	CREMAD $\uparrow$	IEMOCAP $\uparrow$
Wav2Vec2 FT	56.53 [22]	46.02 [23]	69.90 [24]
SELM	75.70	88.16	73.09

Table 2: *In-domain performance of SELM on three datasets. Similar to SELM, the benchmark numbers also use wav2vec2-base embeddings to extract acoustic features.*

### 5.2. Out-of-Domain

In this Section, we compare SELM's Out-of-Domain (OOD) performance against benchmark models from Literature. The existing literature models are also not trained on these three datasets, making this a valid OOD comparison. Table 3 presents the OOD experimental results for RAVDESS, CREMAD and IEMOCAP. The first half of the table presents performance on all emotion classes in the dataset, while the lower half (denoted with \*) presents numbers on a subset of classes. The subset of classes are chosen to be the four primary emotions: *happy, sad, angry, neutral*. The results show SELM outperforms existing models on all three datasets for both all-class and 4-class setup. Moreover, some of the datasets contain emotions that SELM has not seen during training, and hence showcasing the generalization ability of the model.

### 5.3. Few-Shot Learning

In this Section, we compare Few-Shot learning performance of SELM against literature benchmarks. The results are shown in Table 4. The table shows the performance of the model when it is finetuned on either 4 examples per class (4-shot) or 8 examples

Model	Out-of-Domain Dataset		
	RAVD $\uparrow$	CREMA $\uparrow$	IEMOC $\uparrow$
Random	12.50	16.70	-
CLAP [7]	16.00	17.80	13.71
Pengi [6]	20.32	18.46	-
MMS large [25]	13.50	17.20	-
Whisper medium.en [25]	15.30	20.90	-
Whisper medium [25]	16.70	19.90	-
Whisper large-v2 [25]	15.10	20.20	-
AudioFlamingo [12]	20.90	26.50	-
SELM	<b>24.51</b>	<b>28.30</b>	<b>21.42</b>
CLAP* [7]	29.48	31.05	33.72
LanSER PS* [11]	-	15.90	30.90
LanSER CM* [11]	-	23.50	34.30
EmoCLAP* [26]	38.46	35.22	-
SELM*	<b>52.53</b>	<b>42.79</b>	<b>40.02</b>

Table 3: OOD Performance of different models across three datasets. The dataset is considered OOD when labeled or unlabeled audio is not used during training or unsupervised adaptation during testing. The \* symbol indicates only four emotion classes (anger, happiness, sadness, and neutral) are used for evaluation. The metric used is unweighted.

per class (8-shot). For the all-class settings (presented in the earlier rows of the table), we compare SELM’s performance against Audio-Flamingo [12]. The AudioFlamingo reports 4-shot and 8-shot numbers, however the testing strategy and split is not reported. For the 4 class emotion classification, we compare against LanSER [11], which reports the performance when the model is fine-tuned using 10% of the data. We note that 10% data from target domain is significantly more than 4-shot. For example, 10% of data from CREMA-D is equivalent to 744 audio files and hence the setups are not directly comparable. Due to the lack of Few-Shot Learning approaches in SER, we compare SELM against LanSER

The results shown in Table 4 lead to multiple conclusions. First, SELM performs better than prior work on CREMA-D and IEMOCAP. Second, 8-shot always leads to performance improvement over 4 shot. Similarly, 16-shot performs better than 8-shots, i.e. providing SELM more target domain data will lead to better performance. However, there SELM Few-Shot has some limitations. First, the AudioFlamingo performs better on RAVDESS in all-class settings than SELM. We believe that this is due to the specific training strategy of AudioFlamingo which enables use of In-Context Learning. Moreover, AudioFlamingo is trained on 6 million instances, which is 20 times higher than our settings. Second, in RAVDESS, the 4-shot leads to higher performance than 8-shot in the 4-class settings. This might be because RAVDESS has song and speech samples, while SELM has never encountered songs in training data. As songs have different audio distribution from speech samples, the parameter-efficient Few-Shot learning is not sufficient to improve performance and adapt to the song distribution in a few examples.

#### 5.4. Ablation on Few-Shot Learning

To understand which parameter of SELM to update for Few-Shot Learning, we perform multiple ablation studies. The parameters under consideration are: linear layer that is part of the decoder (AL-Dec) or the linear layer part of the encoder (AL-Enc), or the Audio Mapper Transformer (AT) or the Text Mapper Transformer (TT). The results of the ablation study are shown in Figure 2. We find that finetuning Text mapper from SELM leads to the best performance improvement when data distribution is similar. However, for completely different data distribution, adapting a parameter in audio mapper performs better. For example, RAVDESS contains songs which SELM has not seen in training data. To adapt to a widely different audio distribution,

Model	Setup	Dataset for Few-Shot Learning		
		RAVD $\uparrow$	CREMA $\uparrow$	IEMOC $\uparrow$
AudioFlam.	4-shot	-	30.47	-
AudioFlam.	8-shot	<b>35.20</b>	31.80	-
SELM	4-shot	30.09	30.10	25.32
SELM	8-shot	31.81	<b>32.27</b>	<b>27.01</b>
LanSER PS*	10%	-	35.50	42.00
LanSER CM*	10%	-	43.70	50.00
SELM*	4-shot	<b>57.14</b>	43.93	50.17
SELM*	8-shot	55.77	<b>46.02</b>	<b>50.17</b>

Table 4: Few-Shot Learning results of different models and methods on OOD dataset. The dataset is considered OOD when labeled or unlabeled audio is not used during training or unsupervised adaptation during testing. The \* symbol indicates only four emotion classes (anger, happiness, sadness, and neutral) are used for evaluation. The metric used is unweighted accuracy.

requires training parameters in the audio mapper.




Figure 2: Ablation study on the parameters to update for Few-Shot Learning. The first graph and second graph show performance improvement achieved by finetuning different parts of the model under 4-shot and 8-shot settings for the three datasets.

## 6. Limitations

SELM achieves SoTA Out-of-Domain performance. Moreover, by using Few-Shot Learning, the OOD performance can be further improved. However, SELM has limitations. First, SELM is trained only on spoken English datasets and hence its speech emotion recognition will drop on other languages. Second, the Few-Shot Learning improves SELM’s performance on OOD data. However, the improvements are minor when the audio distribution is significantly different.

## 7. Conclusion

In this work, we provide an alternate formulation for Speech Emotion Recognition, posing it as a text token generation task, consisting of both acoustic and language models. We build SELM based on this formulation and test on multiple SER datasets. The results show that SELM generalizes better than the literature models on Out-of-Domain and Few-Shot scenarios. In Out-of-Domain settings, SELM achieves 17% and 7% relative accuracy improvements in RAVDESS and CREMA-D respectively. In Few-Shot setting (4-shot and 8-shot), on average SELM achieves better performance over baseline models.

## 8. References

- [1] H. Dharmyal, B. Raj, and R. Singh, "Positional encoding for capturing modality specific cadence for emotion detection," 2022.
- [2] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in neural information processing systems*, 2020.
- [3] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [4] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *International Conference on Machine Learning*. PMLR, 2023, pp. 28 492–28 518.
- [5] A. Saliba, Y. Li, R. Sanabria, and C. Lai, "Layer-wise analysis of self-supervised acoustic word embeddings: A study on speech emotion recognition," *arXiv preprint arXiv:2402.02617*, 2024.
- [6] S. Deshmukh, B. Elizalde, R. Singh, and H. Wang, "Pengi: An audio language model for audio tasks," *arXiv preprint arXiv:2305.11834*, 2023.
- [7] B. Elizalde, S. Deshmukh, M. A. Ismail, and H. Wang, "Clap learning audio concepts from natural language supervision," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [8] Y. Gong, H. Luo, A. H. Liu, L. Karlinsky, and J. Glass, "Listen, think, and understand," *arXiv preprint arXiv:2305.10790*, 2023.
- [9] Y. Gong, A. H. Liu, H. Luo, L. Karlinsky, and J. Glass, "Joint audio and speech understanding," in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–8.
- [10] S. Deshmukh, B. Elizalde, and H. Wang, "Audio Retrieval with WavText5K and CLAP Training," in *Proc. INTERSPEECH 2023*, 2023, pp. 2948–2952.
- [11] T. Gong, J. Belanich, K. Somandepalli, A. Nagrani, B. Eoff, and B. Jou, "Lanser: Language-model supported speech emotion recognition," *arXiv preprint arXiv:2309.03978*, 2023.
- [12] Z. Kong, A. Goel, R. Badlani, W. Ping, R. Valle, and B. Catanzaro, "Audio flamingo: A novel audio language model with few-shot learning and dialogue abilities," 2024.
- [13] H. Cao, D. G. Cooper, M. K. Keutmann, R. C. Gur, A. Nenkova, and R. Verma, "CREMA-D: Crowd-sourced emotional multimodal actors dataset," *IEEE transactions on affective computing*, vol. 5, no. 4, pp. 377–390, 2014.
- [14] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [15] D. Tompkins, D. Emmanouilidou, S. Deshmukh, and B. Elizalde, "Multi-view learning for speech emotion recognition with categorical emotion, categorical sentiment, and dimensional scores," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [16] R. Lotfian and C. Busso, "Building naturalistic emotionally balanced speech corpus by retrieving emotional speech from existing podcast recordings," *IEEE Transactions on Affective Computing*, vol. 10, no. 4, pp. 471–483, 2017.
- [17] A. Zadeh, R. Zellers, E. Pincus, and L.-P. Morency, "Mosi: multimodal corpus of sentiment intensity and subjectivity analysis in online opinion videos," *arXiv preprint arXiv:1606.06259*, 2016.
- [18] A. B. Zadeh, P. P. Liang, S. Poria, E. Cambria, and L.-P. Morency, "Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2018, pp. 2236–2246.
- [19] S. Poria, D. Hazarika, N. Majumder, G. Naik, E. Cambria, and R. Mihalcea, "Meld: A multimodal multi-party dataset for emotion recognition in conversations," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 527–536.
- [20] S. R. Livingstone and F. A. Russo, "The ryerson audio-visual database of emotional speech and song (ravdess): A dynamic, multimodal set of facial and vocal expressions in north american english," *PLOS ONE*, vol. 13, no. 5, pp. 1–35, 05 2018. [Online]. Available: <https://doi.org/10.1371/journal.pone.0196391>
- [21] 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, pp. 335–359, 2008.
- [22] C. Luna-Jiménez, R. Kleinlein, D. Griol, Z. Callejas, J. M. Montero, and F. Fernández-Martínez, "A proposal for multimodal emotion recognition using aural transformers and action units on ravdess dataset," *Applied Sciences*, vol. 12, no. 1, p. 327, 2021.
- [23] O. C. Phukan, A. B. Buduru, and R. Sharma, "A comparative study of pre-trained speech and audio embeddings for speech emotion recognition," *arXiv preprint arXiv:2304.11472*, 2023.
- [24] L.-W. Chen and A. Rudnicky, "Exploring wav2vec 2.0 fine tuning for improved speech emotion recognition," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [25] R. Ma, A. Liusie, M. J. F. Gales, and K. M. Knill, "Investigating the emergent audio classification ability of asr foundation models," 2023.
- [26] H. Dharmyal, B. Elizalde, S. Deshmukh, H. Wang, B. Raj, and R. Singh, "Describing emotions with acoustic property prompts for speech emotion recognition," *arXiv preprint arXiv:2211.07737*, 2022.