



**HAL**  
open science

# MSP-Podcast SER Challenge 2024: L'antenne du Ventoux Multimodal Self-Supervised Learning for Speech Emotion Recognition

Jarod Duret, Mickael Rouvier, Yannick Estève

► **To cite this version:**

Jarod Duret, Mickael Rouvier, Yannick Estève. MSP-Podcast SER Challenge 2024: L'antenne du Ventoux Multimodal Self-Supervised Learning for Speech Emotion Recognition. Odyssey 2024, Jun 2024, Quebec, France. hal-04628187

**HAL Id: hal-04628187**

**<https://hal.science/hal-04628187v1>**

Submitted on 7 Jul 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# MSP-Podcast SER Challenge 2024: L'antenne du Ventoux

## Multimodal Self-Supervised Learning for Speech Emotion Recognition

Jarod Duret, Mickael Rouvier, Yannick Estève

LIA

Avignon Universite, France

first.last@univ-avignon.fr

### Abstract

In this work, we detail our submission to the 2024 edition of the MSP-Podcast Speech Emotion Recognition (SER) Challenge. This challenge is divided into two distinct tasks: Categorical Emotion Recognition and Emotional Attribute Prediction. We concentrated our efforts on Task 1, which involves the categorical classification of eight emotional states using data from the MSP-Podcast dataset. Our approach employs an ensemble of models, each trained independently and then fused at the score level using a Support Vector Machine (SVM) classifier. The models were trained using various strategies, including Self-Supervised Learning (SSL) fine-tuning across different modalities: speech alone, text alone, and a combined speech and text approach. This joint training methodology aims to enhance the system's ability to accurately classify emotional states. This joint training methodology aims to enhance the system's ability to accurately classify emotional states. Thus, the system obtained F1-macro of 0.35% on development set.

### 1. Introduction

Speech Emotion Recognition (SER) represents a challenging area of research. The complexity arises from the nuanced, subjective nature of emotional expression in speech and the challenge of extracting effective feature representations. Despite these difficulties, understanding emotions is crucial for enhancing human-computer interaction, as emotions significantly influence reasoning, decision-making, and social interactions. In the realm of speech and text, emotions, while subjective, are essential to convey meaning and intent. In recent years, advances in deep learning have contributed to notable improvements in the performance of emotion recognition systems by leveraging highly effective features extracted from deep neural networks. In the pursuit of advancing SER capabilities, the MSP-Podcast SER Challenge 2024[1] provides a fertile ground for exploring novel methodologies and techniques for emotion recognition from naturalistic speech data. Our focus lies primarily on Task 1: Categorical Emotion Recognition, which involves classifying speech segments into eight specified emotional states: Anger (A), Happiness (H), Sadness (S), Fear (F), Surprise (U), Contempt (C), Disgust (D), and Neutral (N).

This paper details our submission to the 2024 edition of the MSP-Podcast SER Challenge. Our approach employs an ensemble of models, each trained independently and then fused at the score level using a Support Vector Machine (SVM) classifier. The models were trained using various strategies, including Self-Supervised Learning (SSL) fine-tuning across different modalities: speech alone, text alone, and a combined speech and text approach. This joint training methodology aims to

enhance the system's ability to accurately classify emotional states.

The paper is structured as follows: In Section 2, we introduce the datasets of MSP-Podcast and the task, while section 3 position our study in regards with some related works. Section 4 presents the general architecture of our system. The components of the sub-systems are described in Section 5. The sub-systems themselves are presented in Section 6, as well as the fusion approach in Section 7. Finally, in Section 8, we present the results and discuss them.

### 2. MSP-Podcast SER Challenge 2024

MSP-Podcast[2] is a large naturalistic speech emotion corpus featuring speech segments sourced from an audio-sharing website. The corpus is annotated for both categorical emotion recognition and emotional attribute prediction. The training set consists of 68,119 speaking turns and the development set contains 19,815 speaking segments. The test set comprises 2,347 unique segments from 187 speakers, with the labels not publicly disclosed. The selection of segments for the test set was carefully curated to ensure a balanced representation across primary categorical emotions.

Two tasks were proposed, we are only participating in the first task. The first task involves categorical classification within eight specified emotional states: Anger (A), Happiness (H), Sadness (S), Fear (F), Surprise (U), Contempt (C), Disgust (D), and Neutral (N). The test set for the challenge has a balanced distribution across the emotional categories. Performance evaluation and model ranking on the leaderboard are based on the Macro-F1 score. The Macro-F1 score is calculated by first computing the F1 score separately for each class, which is the harmonic mean of precision and recall for that class, and then taking the average of these F1 scores.

### 3. Related Work

#### 3.1. Speech Emotion Recognition

SER involves a two-step process: feature extraction and emotion classification. Early SER research focused on hand-crafted features like pitch, energy, and Mel-frequency cepstral coefficients (MFCCs) along with traditional machine learning methods, including Markov models, Gaussian mixture models, and support vector machines for classification [3, 4]. More recently, neural-based models have started to replace traditional machine learning approaches [3][4]. Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) have demonstrated improved performance in emotion recognition tasks [5, 6]. Additionally, transfer learning, particularly

self-supervised pretraining, has gained traction in SER. Models like Wav2Vec2 2.0 [7], WavLM [8] and HuBERT [9] have achieved state-of-the-art results in this domain [10, 11].

### 3.2. Text Emotion Recognition

Text Emotion Detection and Recognition (TEDR) has significantly evolved over the past few years, transitioning from traditional machine learning approaches to deep learning models. Initial methods primarily utilized classifiers such as Support Vector Machine (SVM) and Maximum Entropy (MaxEnt) classifiers, later approaches increasingly relied on deep learning models in combination with different word embedding methods. For example, an emotion detection model that combines Long Short-Term Memory (LSTM) networks with Convolutional Neural Networks (CNN) was introduced [12]. This model integrates various word embeddings, including GloVe [13] and FastText [14], to capture semantic nuances more effectively. More recently, approaches based on transformer pre-trained language models [15] have begun to emerge, offering remarkable breakthroughs in text emotion detection. In [15], the authors conducted a comprehensive comparison of models including BERT [16], RoBERTa [17], DistilBERT [18], and XLNet [19], with RoBERTa emerging as the model achieving the best performance.

## 4. Overview of the approach

The system was developed as a two-level architecture. Given a speech segment, the first level extracts outputs from categorical emotion recognition based on various sub-systems. The outputs of these sub-systems are fed to a SVM. Five different set of sub-systems are used.

We formulate the sub-system as a mapping from the continuous speech domain into the discrete domain of categorical labels of emotion. As depicted in Figure 1, to achieve this, we first use an encoder (speech and/or text encoder). These encoders have been trained on unlabeled data and are capable of extracting highly robust feature representations. Following the encoder stage, we employ a pooling strategy to aggregate these features over time, ensuring a fixed-length representation regardless of the original speech duration. This representation then feeds into a classifier layer, which serves as the final component of our architecture. This layer is responsible for mapping the aggregated features to the desired categorical labels of emotion, thus completing the process of transforming continuous speech into discrete emotion predictions.

## 5. Sub-System components

In this section, we describe the various components used in the sub-systems. In Section 5.1, we describe the various speech and text encoders used in our study. Following that, Section 5.2 is dedicated to detailing the pooling techniques employed. The discussion continues in Section 5.3, where we delve into the classifier used. Finally, Section 5.4 covers the aspects of speech audio data augmentation.

### 5.1. Encoder

In the realm of speech and text processing, the choice of encoders is essential for provide powerful deep feature learning. Indeed, these encoders are trained on large unannotated datasets, enabling them to learn rich, complex patterns without

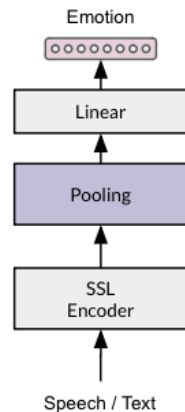


Figure 1: Illustration of our speech emotion recognition system.

the need for manually labeled data.

For speech encoder, we leverage the capabilities of WavLM<sup>1</sup>, a state-of-the-art Self-Supervised Learning (SSL) speech model designed to discover speech representation that encode pseudo-phonetic or phonotactic information. WavLM is distinguished by its robustness in handling a wide range of audio scenarios, including noisy environments. Additionally, we also used HuBERT<sup>2</sup>, another leading speech encoder model. For both model, features are extracted, at the acoustic frame-level i.e. for short speech segments of 20 ms duration.

For text encoder, we employ RoBERTa<sup>3</sup>, which is builds on the BERT architecture but optimizes its pre-training process for more efficient learning and better performance across a variety of Natural Language Processing (NLP) tasks. Features are extracted for every words.

### 5.2. Pooling

Given that speech segments vary in length, we use pooling to aggregate the features given by the encoder across time, ensuring a fixed-length representation regardless of the original speech duration. Two different types of pooling have been used:

- **mean-pooling** achieves this by averaging the feature values given by the encoder over the time dimension, which effectively summarizes the overall characteristics of the speech signal into a single unified representation.
- **attention-pooling** unlike mean-pooling leverages a weighted average, where the weights are learned through the model [20]. This allows the model to focus on more relevant parts of the speech signal, potentially capturing nuanced dynamics that mean pooling might overlook.

### 5.3. Classifier

For the classifier component, we assume that the SSL encoders have successfully captured all necessary information for predicting the targeted categorical emotion label. We choose a simple yet efficient approach by integrating a linear layer to function as the classifier. This linear classifier takes the pooled feature vectors and assigns them to the a given emotion label.

<sup>1</sup><https://huggingface.co/microsoft/wavlm-large>

<sup>2</sup><https://huggingface.co/facebook/hubert-large-ll60k>

<sup>3</sup><https://huggingface.co/FacebookAI/roberta-base>

## 5.4. Data Augmentation

The MSP-Podcast corpus includes, for each segment, the emotion label, text transcription, speaker ID, and gender. However, the test set is provided as speech only, without any annotations. To utilize a text encoder, we needed to automatically generate transcriptions for the test set. For this purpose, we employed the Whisper [21] speech recognition model. Additionally, to ensure consistency and avoid any discrepancies between provided transcriptions and Whisper-generated transcriptions, we computed transcriptions for the entire dataset.

We observed a significant discrepancies in the distribution of emotion classes within both the Training and Development sets. Additionally, a considerable number of samples are labeled as "X," indicating that no consensus was reached for these samples. To minimize the mismatch, we decided to automatically recompute the consensus for samples labeled as "X". For each sample, we have access to the labels provided by each annotator and the consensus. The consensus is determined by identifying the most frequently associated label. In scenarios where no single label predominates due to an even distribution of votes among the labels, the label "X" is assigned to indicate the lack of a clear consensus.

Our method for recomputing the consensus is detailed as follows: First, we calculate a score for each evaluator by determining the ratio of their annotations that match the consensus to their total number of annotations. Then, we discard all annotations from evaluators whose score falls below a predefined threshold  $K$ . Finally, we recompute the consensus for the whole dataset. In scenarios where there is a tie between a specific label and the Neutral ("N") label, we opt to drop the label "N". Samples with newly attributed labels from the Training and Development sets were added to the training set.

## 6. Sub-Systems

In this section, we provide a detailed description of the five distinct sub-systems employed in the fusion process.

### 6.1. Sub-System A : WavLM

This system is based on a WavLM, a mean pooling strategy and the output is a softmax loss function. In order to optimize this architecture, we employ the Adam optimizer. We fine-tune the pre-trained WavLM model during the training phase. This fine-tuning allows the WavLM model to adjust its parameters specifically towards recognizing emotions in speech, leveraging the rich, pre-learned representations and tailoring them to our domain of interest. We set the mini-batch size to 16 and 10 steps. And no data-augmentation is done on speech segment.

### 6.2. Sub-System B : Jeffreys Loss

This system is identical to System A; however, we propose in this system to replace the softmax loss function by Jeffreys loss function. The Jeffreys Loss is given in the Equation 1. Incorporating this divergence into the cross-entropy loss function enables the maximization of the target value of the output distribution while smoothing the non-target values.

$$\mathcal{L} = -\log(p_k) - \alpha \frac{\sum_{i \neq k} \log p_i}{K-1} + \beta \frac{\sum_{i \neq k} p_i \log p_i}{1-p_k} \quad (1)$$

### 6.3. Sub-System C : Joint Wav2vec2-WavLM

This system is based on a joint SSL audio strategy as depicted in Figure 2, wherein the upstream model is performed by jointly training a WavLM and Hubert SSL model. The input speech is fed into both SSL models. Fine-tuning of the upstream model is achieved through simultaneous training alongside a straightforward network that implements mean pooling, leading to a Linear Classifier, as illustrated in Figure 2.

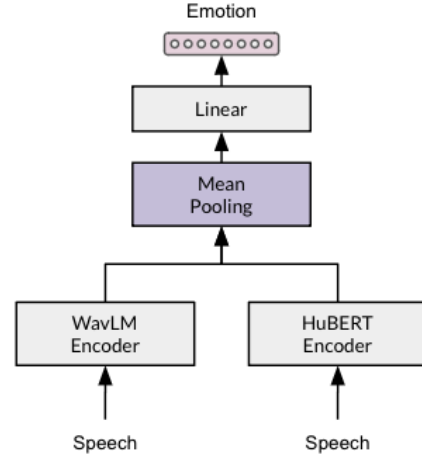


Figure 2: Illustration of the dual speech encoder emotion recognition system.

### 6.4. Sub-System D : WavLM and RoBERTa

The following describes the SER model depicted in Figure 3. The WavLM encoder is divided into two parts: the CNN feature extractor and the transformer-based encoder. We chose to freeze the CNN feature extractor part, fixing all the parameters of these CNN layers and only fine-tune the parameters of transformer blocks. This method of partial fine-tuning acts as a strategy for domain adaptation. It is designed to maintain the integral feature extraction capabilities of the lower layers, thereby enabling the model to adjust to new tasks efficiently without compromising its overall performance. For the text encoder, we opt to fine-tune all parameters of the RoBERTa model. During the fine-tuning process, we apply three different schedulers to adjust the fine-tuning learning rates of the WavLM and RoBERTa encoders, as well as the learning rate of the classifier model. Each scheduler utilizes the Adam Optimizer in conjunction with the NewBob technique, which anneals the learning rates based on the performance during the validation stage. The fine-tuning learning rates for both the WavLM and RoBERTa encoders are set to  $10^{-5}$ , while the learning rate for the classifier model is set to  $10^{-4}$ . This model is trained using negative log-likelihood loss (NLL).

### 6.5. Sub-System E : Data Augmentation

The system E is identical to the system D described in section 6.4. However, this system incorporates the data augmentation technique outlined in Section 5.4, setting  $K = 50\%$ . This adjustment aims to mitigate the notable imbalances observed in the distribution of emotion classes.

Table 1: Performance comparison of individual sub-systems and their fusion across all emotion classes, with overall accuracy and F1-Macro scores.

	Sub-System A	Sub-System B	Sub-System C	Sub-System D	Sub-System E	Fusion
Anger (A)	0.546	0.552	0.524	0.487	0.490	0.498
Contempt (C)	0.141	0.104	0.164	0.198	0.233	0.271
Disgust (D)	0.118	0.166	0.164	0.191	0.161	0.130
Fear (F)	0.033	0.026	0.055	0.064	0.027	0.027
Happiness (H)	0.589	0.577	0.551	0.591	0.614	0.624
Neutral (N)	0.608	0.604	0.595	0.536	0.606	0.632
Sadness (S)	0.315	0.360	0.318	0.381	0.386	0.395
Surprise (U)	0.202	0.186	0.174	0.211	0.169	0.222
Accuracy	0.509	0.508	0.484	0.466	0.505	0.522
F1-Macro	0.319	0.322	0.318	0.332	0.336	0.350

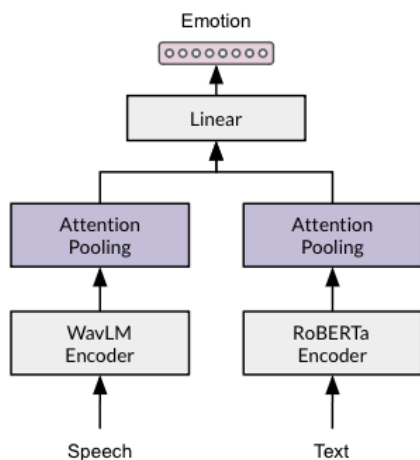


Figure 3: Illustration of our joint speech and text emotion recognition system.

## 7. Fusion systems

The fusion process integrates the outputs from all five sub-systems (A, B, C, D, E) by concatenating them to form a single feature vector. This concatenated vector serves as the input to an SVM classifier, which is trained to predict the emotion class. This fusion approach leverages the diverse strengths of each sub-system, aiming to enhance the overall performance and robustness of emotion classification.

## 8. Experiments

In this section, we describe our experimental setup, analyze results from individual sub-systems and their combined fusion.

### 8.1. Experimental setup

To conduct our experiments, we employed the SpeechBrain toolkit [22], which is built on PyTorch and is specifically designed for speech-processing tasks. Additionally, we utilized the Hugging Face versions of the WavLM, HuBERT, and RoBERTa models. The source code is available on GitHub<sup>4</sup>.

<sup>4</sup>github.com/Chaanks/MSP-Podcast-SER-Challenge-2024

Table 2: Distribution of emotion classes before and after data augmentation.

Emotion Class	Train	Train Augmented
N	25106	25106
H	13440	13440
S	3882	6067
A	3053	4753
U	2897	4328
C	2443	2897
D	1426	2352
F	1139	1681
Total	53386	60624

### 8.2. Results

**Impact of the Speech encoder:** A first experiments aimed at identifying the most effective Speech encoder. Various speech encoders, including Wav2Vec, Hubert and WavLM were evaluated. These experiments were carried out on sub-system D. Table 3, give the results of the experiments, from which we observe that the best results are obtained with the WavLM.

**Impact of the Text encoder:** An additional experiment was conducted using text as input. This experiment, as shown in Table 3, reveals that the performance of the RoBERTa text encoder, achieving a Macro-F1 score of 0.27, is not too far from the best speech encoder, which has a Macro-F1 score of 0.32. This result encourages the exploration of a joint speech and text emotion recognition system.

Table 3: Performance Comparison of Speech and Text Encoders. The table shows Micro-F1 and Macro-F1 scores for speech encoders (Wav2Vec2, WavLM, HuBERT) and a text encoder (RoBERTa).

Model	Micro-F1 $\uparrow$	Macro-F1 $\uparrow$
Wav2Vec2 Large	0.45	0.29
WavLM Large	0.51	0.32
HuBERT Large	0.50	0.31
RoBERTa Base	0.44	0.27

**Impact of the Fusion:** Table 1 shows the results achieved by the different sub-systems. We observe that all the sub-systems achieved F1-Macro scores between 0.32 and 0.34. The best sub-system is sub-system E. And we observe that fusion system led to an improvement, achieving an F1-Macro score of 0.35%.

**Confusion matrix:** Figure 4 shows the confusion matrix obtained from fusion system. Regarding the diagonal elements, we can observe that the fusion model is particularly effective at correctly predicting Neutral (N) and Happiness (H) classes. However, the fusion system struggles more with accurately predicting the Disgust (D) and Fear (F) classes. As for the off-diagonal elements, we can note that there is a tendency for all classes to be misclassified as Neutral (N) and Disgust (D) misclassified as Contempt (C).

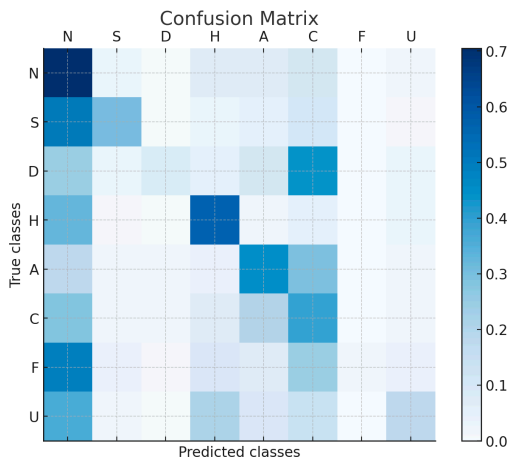


Figure 4: The confusion matrix provided by the fusion system.

## 9. Conclusion

This paper describes LIA’s participation in the MSP-Podcast SER Challenge. Our approach involves developing sub-systems, each of which functions as an emotion classifier. These sub-systems model speech segments using different components to provide varied perspectives. For the final fusion step, we concatenate the outputs of the sub-systems and train a SVM for fusion. The fusion system achieved an F1-Macro score of 0.35% on the development set.

## 10. Acknowledgements

This work was granted access to the HPC resources of IDRIS under the allocation AD011013257R2 made by GENCI.

## 11. References

- [1] L. Goncalves, A. N. Salman, A. Reddy Naini, L. Moro-Velazquez, T. Thebaud, L.P. Garcia, N. Dehak, B. Sisman, and C. Busso, “Odyssey2024 - speech emotion recognition challenge: Dataset, baseline framework, and results,” in *Odyssey 2024: The Speaker and Language Recognition Workshop*, Quebec, Canada, June 2024, vol. To appear.
- [2] 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, October-December 2019.
- [3] B. Schuller, S. Steidl, and A. Batliner, “The INTER-SPEECH 2009 emotion challenge,” in *Interspeech 2009*, 2009, ISCA.
- [4] B. Schuller, S. Steidl, and A. et al. Batliner, “The INTER-SPEECH 2013 computational paralinguistics challenge: social signals, conflict, emotion, autism,” in *Interspeech 2013*, 2013, ISCA.
- [5] L. Yuanchao, Z. Tianyu, and K. Tatsuya, “Improved end-to-end speech emotion recognition using self attention mechanism and multitask learning,” in *Interspeech 2019*, 2019, ISCA.
- [6] C. Etienne, G. Fidanza, A. Petrovskii, L. Devillers, and B. Schmauch, “CNLSTM architecture for speech emotion recognition with data augmentation,” in *Workshop on Speech, Music and Mind (SMM 2018)*, 2018, ISCA.
- [7] A. Baeviski, 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*, vol. 33, pp. 12449–12460, 2020.
- [8] Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, et al., “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [9] K. Lakhotia, E. Kharitonov, W. Hsu, Y. Adi, A. Polyak, B. Bolte, T. Nguyen, J. Copet, A. Baeviski, and A. Mohamed, “On generative spoken language modeling from raw audio,” *Transactions of the Association for Computational Linguistics*, vol. 9, pp. 1336–1354, 2021.
- [10] M. Macary, M. Tahon, Y. Estève, and A. Rousseau, “On the use of self-supervised pre-trained acoustic and linguistic features for continuous speech emotion recognition,” in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 373–380.
- [11] Y. Wang, A. Boumadane, and A. Heba, “A fine-tuned wav2vec 2.0/hubert benchmark for speech emotion recognition, speaker verification and spoken language understanding,” *arXiv preprint arXiv:2111.02735*, 2021.
- [12] Marco Polignano, Pierpaolo Basile, Marco de Gemmis, and Giovanni Semeraro, “A comparison of word-embeddings in emotion detection from text using bilstm, cnn and self-attention,” in *Adjunct publication of the 27th conference on user modeling, adaptation and personalization*, 2019, pp. 63–68.
- [13] Jeffrey Pennington, Richard Socher, and Christopher D Manning, “Glove: Global vectors for word representation,” in *Proceedings of the 2014 conference on empirical*

*methods in natural language processing (EMNLP)*, 2014, pp. 1532–1543.

- [14] Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, H erve J egou, and Tomas Mikolov, “Fast-text. zip: Compressing text classification models,” *arXiv preprint arXiv:1612.03651*, 2016.
- [15] Acheampong Francisca Adoma, Nunoo-Mensah Henry, and Wenyu Chen, “Comparative analyses of bert, roberta, distilbert, and xlnet for text-based emotion recognition,” in *2020 17th International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP)*. IEEE, 2020, pp. 117–121.
- [16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [17] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov, “Roberta: A robustly optimized bert pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.
- [18] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf, “Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter,” *arXiv preprint arXiv:1910.01108*, 2019.
- [19] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le, “Xlnet: Generalized autoregressive pretraining for language understanding,” *Advances in neural information processing systems*, vol. 32, 2019.
- [20] Pooyan Safari, Miquel India, and Javier Hernando, “Self-attention encoding and pooling for speaker recognition,” 2020.
- [21] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever, “Robust speech recognition via large-scale weak supervision,” 2022.
- [22] Mirco Ravanelli, Titouan Parcollet, Peter Plantinga, Aku Rouhe, Samuele Cornell, Loren Lugosch, Cem Subakan, Nauman Dawalatabad, Abdelwahab Heba, Jianyuan Zhong, Ju-Chieh Chou, Sung-Lin Yeh, Szu-Wei Fu, Chien-Feng Liao, Elena Rastorgueva, Fran ois Grondin, William Aris, Hwidong Na, Yan Gao, Renato De Mori, and Yoshua Bengio, “SpeechBrain: A general-purpose speech toolkit,” 2021, arXiv:2106.04624.