



What Does it Take to Generalize SER Model Across Datasets? A Comprehensive Benchmark

Adham Ibrahim¹, Shady Shehata^{1,2}, Ajinkya Kulkarni^{1,3}, Mukhtar Mohamed¹, Muhammad Abdul-Mageed^{1,2,4}

¹Mohamed bin Zayed University of Artificial Intelligence, UAE; ²Invertible AI, UAE; ³IDIAP, Switzerland; ⁴University of British Columbia, Canada

{adham.ibrahim, shady.shehata}@mbzuai.ac.ae

Abstract

Speech emotion recognition (SER) is essential for enhancing human-computer interaction in speech-based applications. Despite improvements in specific emotional datasets, there is still a research gap in SER's capability to generalize across real-world situations. In this paper, we investigate approaches to generalize the SER system across different emotion datasets. In particular, incorporate 11 emotional speech datasets and illustrate a comprehensive benchmark on the SER task. We also address the challenge of imbalanced data distribution using over-sampling methods when combining SER datasets for training. Furthermore, we explore various evaluation protocols for adeptness in the generalization of SER. Building on this, we explore the potential of Whisper for SER, emphasizing the importance of thorough evaluation. Our approach is designed to advance SER technology by integrating speaker-independent methods.

Index Terms: speech emotion recognition, human-computer interaction, computational paralinguistics

1. Introduction

Emotions are intrinsic characteristics of human communications expressed through speech modality and plays crucial role in human-machine interfaces [1]. Over the past decades, research has been focused on various feature extraction methods along with machine learning. This includes traditional approaches like Mel-frequency cepstral coefficients (MFCCs)[2], linear prediction cepstral coefficients (LPCCs)[3], and prosodic features[4]. Deep learning techniques such as deep neural network DNN[5], recurrent neural network (RNN)[6], and convolutional neural networks (CNN)[7] have also gained prominence. The last three years have seen a considerable exploration of HuBERT[8] and Wav2Vec2[9] models in SER [10, 9], demonstrating their potential for robust emotion recognition.

IEMOCAP [11] is considered the benchmark dataset for discrete speech emotion recognition. An extensive evaluation of SER systems has been conducted using IEMOCAP, employing various machine learning methods such as support vector machines (SVM), LSTM, CNN, and ensemble learning [12, 13]. However, the use of Self-supervised learning (SSL) speech models (e.g., Hubert [8] and Wav2Vec2 [9]) has emerged as the state-of-the-art approach in SER, producing the best performance on the IEMOCAP dataset [9]. Another very important acted dataset is (RAVDESS) [14]. RAVDESS was also involved in benchmarking many different SER systems [15, 16]. A dataset closely similar to RAVDESS is CREMA-D [17], which comprises 12 distinct spoken sentences performed by 91 different actors, resulting in a total of 7442 unique audio files.

Among the more compact SER databases available, the Surrey Audio-Visual Expressed Emotion (SAVEE) [18] and the

Toronto Emotional Speech Set (TESS) [19] stand out. Both datasets have only been benchmarked using traditional machine learning techniques (SVM, LSTM, CNN) [20, 21, 22]. To the best of our knowledge, these datasets have yet to be tested with newer transformer-based models. MELD dataset has a unique conversational setup of 13,000 utterances from 1,433 dialogues from the TV series Friends. This compilation aims to address the complexities of emotion recognition within conversational contexts, a task known for its challenges. Other SER datasets include ASVP-ESD[23], EmoV-DB[24], EmoFilm[25], JL-Corpus[26], and ESD[27].

Despite the availability of various datasets for speech emotion recognition, there's a consensus that SER still trails behind other tasks [28, 29, 30]. This shortfall is attributed to the absence of a large, universal dataset capable of bridging the significant disparities among existing datasets and facilitating adaptation to real-world scenarios. A recent study [16] investigated the generalization capabilities and the possibility of merging different SER datasets by applying a set of cross-validation experiments, considering both single datasets and combinations of them, on RAVDESS, TESS, CREMA-D, and IEMOCAP. Their results confirm that SER models do not generalize well across datasets (training and testing on different datasets) and that merging datasets can mitigate this problem, as SER performance improves with access to a larger and more varied collection of data points.

A key challenge in speech emotion recognition stems from the substantial variance among available datasets, attributed to differences in setup (acted, natural, and elicited), recording quality, and subjective emotion perception by speakers and annotators. Consequently, most systems struggle to generalize across datasets, hindering performance in real-world applications [16]. Moreover, emotional states influence spoken sentences, compounded by personal attributes like age and gender, posing obstacles for SER generalization and hindering the development of speech-independent systems, particularly in tasks like speech emotion recognition (SER). This challenge is underscored by human behavior, as demonstrated in an experiment by Schuller et al. [31] where the accuracy of 12 participants in recognizing expressed emotions decreased from 87.3% with familiar speakers to 64.7% when identifying emotions from speakers they had not previously encountered.

Due to variations in emotion classes, emotion labeling, amount of emotional data per dataset, SER more arduous task to generalize. In this paper, we present novel comprehensive benchmarking results across 11 SER datasets and various experiments showcasing the vital aspects of generalization and performance improvement of SER.




Figure 1: Whisper based speech emotion recognition system

2. Methodology

In this section, we highlight our primary contributions: Our SER system based on Whisper, the combination of different datasets, and evaluation strategy.

2.1. Model Architecture

We illustrate the SER model architecture, which we use throughout this work for various experiments across various SER datasets. This architecture aligns with the current state-of-the-art (SoTA) and is representative of the current best practices in speech emotion recognition [10].

Our SER first uses Whisper-based feature extraction [32] to map speech into latent representation through a transformer-based encoder-decoder network as shown in Fig 1. We use the base version of Whisper utilizing both the encoder and decoder to extract fixed-size embedding and fine-tune the whole model with the rest of the network. Thereafter, this latent representation as features is given to feed-forward of fully connected layers: starting from a size of 4096, it sequentially narrows down through layers of 2048, 1024, and 512, with each layer followed by a ReLU activation function to add non-linearity. This setup culminates in an output layer tailored to the number of target classes, making it suitable for classification purposes. To our knowledge, Whisper-based SER has not been explored for ASR tasks. We implement a basic five-layer feed-forward neural network, where the last layer projects into the different emotion labels. Cross-entropy serves as our loss function, and we maintain a constant learning rate of 0.00001. We train all variants of SER on NVIDIA A100-SXM GPUs equipped with 40GB of memory.

2.2. Datasets Preparation

Emotional databases in speech emotion recognition are classified into acted, elicited, and natural categories, each with unique characteristics and challenges. Acted databases consist of recordings from trained actors simulating specific emotions, offering control but relying on the actors’ skills[33]. Elicited databases capture genuine emotional responses under designed scenarios requiring ethical consent, while natural databases compile authentic interactions from everyday life, presenting a wide authenticity spectrum but limited emotional range[34].

Our study conducts a detailed analysis of individual and combined datasets to understand the impact of similar emotions on model accuracy and generalization on 11 datasets in the SER domain: **IEMOCAP**[11]: Approximately 12 hours from 10 speakers, with scripted and improvised dialogues across seven emotion classes (anger, happiness, sadness, excited, frustrated, and neutral). **MELD**[35]: 13,708 utterances over 12 hours from the Friends TV series, categorized into seven emotions (anger, disgust, sadness, joy, neutral, surprise, fear) with 10 main speakers and around 290 secondary speakers. **ASVP-ESD**[23]: A multi-lingual corpus from movies, YouTube, and real interactions, focusing on English speech with seven emo-

Table 1: Distribution of emotion samples across 11 SER datasets

Emotion	Dataset Name											
	ASVP-ESD	CREMA-D	EmoF	ESD	IEMOCAP	JL	MELD	TESS	emov.db	savee	ravdess	Total
angry	585	1271	77	3500	1103	240	1607	400	1268	60	192	10303
disgust	612	1271	-	-	-	-	301	400	522	45	184	3335
anxious	-	-	-	-	-	240	-	-	-	-	-	240
apologetic	-	-	-	-	-	240	-	-	-	-	-	240
assertive	-	-	-	-	-	240	-	-	-	-	-	240
concerned	-	-	-	-	-	240	-	-	-	-	-	240
encouraging	-	-	-	-	-	240	-	-	-	-	-	240
excited	-	-	-	-	1041	240	-	-	-	-	-	1281
frustrated	-	-	-	-	1849	-	-	-	-	-	-	1849
fear	544	1271	72	-	-	-	358	400	-	60	192	2897
happy	652	1271	70	3500	595	240	2307	400	-	60	192	9287
neutral	714	1087	-	3500	1708	240	6432	400	1568	120	96	15865
pain	558	-	-	-	-	-	-	-	-	-	-	558
sad	-	1271	74	3500	1084	240	1002	400	-	60	192	7823
surprised	695	-	-	3500	107	-	1403	400	-	60	192	6357
contempt	-	-	50	-	-	-	-	-	-	-	-	50
amused	-	-	-	-	-	-	-	-	1317	-	-	1317
sleepy	-	-	-	-	-	-	-	-	1721	-	-	1721
calm	-	-	-	-	-	-	-	-	-	192	192	192

tions, totaling 6.5 hours from main and secondary speakers. **EmoV-DB**[24]: Covers five emotions with five speakers (four in English, one in French), totaling 9.5 hours. **TESS**[19]: Features two actresses, covering seven emotion classes over 2800 files in 1.6 hours. **EmoFilm**[25]: Samples in English, Italian, and Spanish from 43 movies, focusing on five emotions over 20 minutes. **SAVEE**[18]: British-English corpus with 480 utterances from four speakers, spanning seven emotions in 30 minutes. **RAVDESS**[14]: Focuses on the speech part with 1440 utterances from 24 actors, covering eight emotions over 1.5 hours. **CREMA-D**[17]: Audio-visual dataset with 7442 stimuli from 91 actors, spanning six emotions over 5.3 hours. **JL-corpus**[26]: Contains 2400 sentences from four speakers, covering 10 emotions over 1.4 hours. **ESD**[27]: Offers 350 parallel utterances in five emotion categories from 10 native English speakers, with over 29 hours recorded.

In our investigation, we specifically extracted English speech and corresponding labels from mixed-language datasets like ESD[27], ASVP-ESD[23], and EmoFilm[25]. It’s essential to highlight that the other datasets employed in our study are exclusively in English, ensuring a consistent linguistic foundation for our analysis. Regarding the IEMOCAP dataset, our dataset preparation followed the guidelines outlined in a previous study[36], where emotions were categorized into anger, happiness, sadness, and neutral, with excitement categorized under happiness. We retained the original emotion classes for the remaining datasets. Additionally, we downsampled speech across all speech utterances to 16 kHz. Speaker identification for each audio clip in every dataset was accomplished using the information provided in the dataset metadata.

2.3. Combining datasets for the SER task

A major challenge in speech emotion recognition lies in the considerable variance among the available datasets. This variance stems from differences in their setup (acted, natural, and elicited), the quality of recordings, and the subjective perception of emotions by different speakers and annotators. Consequently, most speech emotion recognition systems do not generalize well across different datasets (training and testing on different datasets) [16] and, by extension, struggle to perform well in real-world applications.

In an effort to address this challenge, [16] showed promising results by combining four of the most famous SER datasets (RAVDESS[14], TESS[19], CREMA-D[17], and

IEMOCAP[11]). In this work, in addition to these four datasets, we added seven more (MELD[35], ASVP-ESD[23], EmoV-DB[24], EmoFilm[25], SAVEE[18], JL-Corpus[26], and ESD[27]) for a total of eleven SER datasets. To the best of our knowledge, combining 11 datasets presents the largest training dataset for the SER task. This study aims to significantly boost performance (against individual datasets) on the SER task, attributable to the considerably larger and more diverse array of training data points.

2.4. Evaluation protocol

Throughout this work, we opted for the leave-one-speaker-out (LOSO) method to evaluate our models. We used accuracy as a way to measure the performance of SER systems. When dealing with datasets where all speakers contributed the same amount of speech, we randomly selected one speaker from each dataset as our test case. Conversely, for datasets where speakers had varying amounts of speech time, such as MELD, ASVP-ESD, and IEMOCAP, we excluded the speaker with the most speech duration. At the end of this process, we ensured that a speaker from every dataset was chosen, guaranteeing a wide-ranging and inclusive evaluation.

3. Experimental Study

In our work, we consider three different sets of experiments for training and evaluation: **1.** We establish an initial benchmark by training each dataset individually using all their original emotion classes. **2.** We use only four emotions (neutral, angry, happy, and sad) to train and test our SER models. We do this for both individual and combined settings. **3.** We use five emotions adding surprise to the previous four to train and test our SER models. We do this for both individual and combined settings.

We used a consistent training setup involving 5-fold cross-validation. This means we divided the data so that each fold used a different speaker for validation and the rest for training. For datasets with less than five speakers, we adjusted the number of folds to match the number of speakers. Additionally, with the TESS dataset, which has only two speakers, we used one for testing and the other for training, applying 5-fold cross-validation on the training speaker’s data. In the case of the EmoFilm dataset, we used 85% of the data for training, still following the 5-fold cross-validation method. For combined datasets, however, we adapted 5-fold cross-validation based on percentage splits across the dataset to ensure even distribution. It is important to note that we leave one speaker from each dataset as a test set right from the start, calculating a weighted average on the speaker-out test set across the 5-folds to ensure uniform and comprehensive evaluation metrics across all experiments.

3.1. Individual and combined 4-emotion dataset evaluation

For every dataset, we train a Whisper-based model using the four primary emotions: neutral, angry, happy, and sad. If a dataset lacks one of these emotions, we include from the dataset whichever available subset from these four emotions. **Table 1** shows that we considered the most common 4 emotions across all datasets. For the combined dataset, we combined 11 SER datasets into a single comprehensive dataset based on the same 4 emotions. We included audio clips with duration ranging from 2 to 13 seconds, to ensure consistency and comparability across the combined dataset. Speech clips within this duration range are the most common duration of data across all datasets. After

Table 2: Comparison of emotion distribution in combined 11 SER datasets with different sampling techniques.

Emotion/Method	4-Emotions				5-Emotions			
	Original	Undersamp	SMOTE	ADASYN	Original	Undersamp	SMOTE	ADASYN
Neutral (neu)	10044	6157	10044	10044	10044	4056	10044	10044
Angry (ang)	7427	6157	10044	10044	7427	4056	10044	10044
Happy (hap)	6286	6157	10044	10044	6286	4056	10044	10044
Sad (sad)	6157	6157	10044	10044	6157	4056	10044	10044
Surprised (sur)	-	-	-	-	4056	4056	10044	10044

combining emotions from various datasets, the data distribution exhibited imbalance, as illustrated in **Table 2**. This imbalance in the distribution of emotional categories across the combined dataset could potentially influence both the training and evaluation phases of our model. Explored four techniques to determine if any could improve model generalization: no sampling, randomly downsampling to the lowest category, and oversampling performed using the Synthetic Minority Over-sampling Technique (SMOTE) [37] and the Adaptive Synthetic Sampling Method for Imbalanced Data (ADASYN)[38]. We applied the SMOTE and ADASYN algorithms to the audio data. The oversampling process was conducted to balance the distribution of emotion labels, ensuring that low-frequency emotions were over-sampled to match the frequency of the highest emotion category.

3.2. Individual and combined 5-emotions dataset evaluation

This step was taken to ascertain the model’s robustness in handling an expanded set of emotion labels. we trained a Whisper-based model on a set of 5 emotions neutral, angry, happy, surprised, and sad. We train the model on each dataset using the same settings as in the individual training described in the previous experiment in Section 3.1. We repeat training and test settings in combined dataset training in the previous experiment, but with 5 emotions neutral, angry, happy, surprised, and sad.

4. Results and Discussion

In this section, we provide a comprehensive overview of speech emotion recognition performance across different datasets, emotion categories, training criteria, and data sampling techniques, offering valuable insights for the generalization of speech emotion recognition. **Table 3.** illustrates the performance of the Whisper-based speech emotion recognition system developed under various combinations of dataset, and oversampling techniques. In general, the performance of SER varies across datasets and emotion categories. For instance, datasets like `emov_db` consistently demonstrate high accuracy, while others like MELD exhibit lower performance under all conditions. Overall, using original data distribution or using SMOTE sampling tends to yield better performance compared to downsampling and ADASYN sampling. For instance, mean accuracy for training SER using a combination of 4 emotions demonstrates that using original data distribution and SMOTE sampling provides better results than other conditions. This suggests that maintaining the original sample distribution or using synthetic samples for minority classes can enhance model performance. Furthermore, training systems separately on each dataset for 4 emotion and 5 emotion experiments under performed in comparison to combining datasets. Thereafter, performance varies across datasets and emotion categories. For instance, datasets like `emov_db` consistently showcase high accuracy, while others like MELD display comparatively lower performance. The choice of data sampling technique signifi-




Figure 2: *t*-SNE visualization of Whisper model’s embeddings showcasing clusters of emotional speech from various datasets.

Table 3: *Experiment Results for 4-Emotions, 5-Emotions, and N-Emotions, detailing performance across training criteria and data sampling techniques; Tr Co refers to training combination of datasets, Tr. sep. refers to training datasets separate, WA as weighted average, DS as downsampling, UN as original samples, SM refers to SMOTE sampling and AD as ADASYN oversampling*

Dataset Name	4-Emotions WA Test (%)					5-Emotions WA Test (%)					N-Emotions WA Test (%)			
	Emo. No.	Tr. Co.				Tr. Sep.	Emo. No.	Tr. Co.				Tr. Sep.	Emo. No.	Tr. Sep.
		DS	UN	SM	AD			DS	UN	SM	AD			
ESD	4	79.11	77.5	72.73	74.45	76.85	5	78.81	74.45	74.72	74.72	72.04	5	72.04
MELD	4	51.20	54.72	56.11	56.65	63	5	47.66	54.72	47.68	47.68	56.87	7	50
IEMOCAP	4	67.75	70.00	68.79	68.96	72.43	5	66.37	70.00	67.42	67.42	69.5	7	56.1
CREMA-D	4	93.50	98.00	95	93.5	90.41	4	91	98.00	95.50	95.50	90.41	6	80
EMOV.DB	2	99.45	99.64	99.7	99.68	94.4	2	99.6	99.64	99.66	99.66	94.4	5	80.9
ASVP-ESD	3	79.11	77.75	81.25	78.75	84.65	4	73.37	77.75	76.30	76.30	82.35	7	73.33
TESS	4	77.35	71.08	74.9	74.07	66.95	5	62.5	71.08	55.83	55.83	56.62	7	55.52
JL_corpus	4	67.50	56.45	68.95	63.42	53.47	4	53.42	56.45	64.34	64.34	53.47	10	32.25
RAVDESS	4	85.00	87.50	87.50	81.67	80.83	5	87.5	87.50	88.13	88.13	80.5	8	74.69
SAVEE	4	71.67	72.50	77.78	77.78	68.05	5	67.5	72.50	70.45	69.77	73.11	7	67.71
EmoFilm	3	67.50	64.17	62.5	58.53	82.2	3	68.33	64.17	57.5	60.83	82.2	5	53.37
Mean	-	76	78.64	77.24	75.22	75.84	-	72.36	75.11	72.56	73.38	73.74	-	63.20

cantly influences emotion recognition performance. Techniques like SMOTE and ADASYN, which address the class imbalance, result in improved accuracy compared to simple downsampling. This highlights the importance of addressing class imbalance in training data to improve the robustness of emotion recognition models.

Each dataset exhibits unique characteristics concerning the designing process of the dataset such as emotion label annotation, recording conditions of emotions, speaker demographics, and linguistic variations. etc, have an impact on emotion recognition performance. Hence, the factors such as data quality, diversity, and class distribution influence the effectiveness of trained models. For example, datasets like CREMA-D and emov_db, which offer high-quality and diverse speech samples, tend to achieve better performance across different emotion categories. Therefore, it is vital to understand how different training conditions affect model generalization is essential for developing robust speech emotion recognition systems. In conducting our research on 11 SER datasets with Whisper-based models, we encountered several limitations.

Initially, the inherent diversity and imbalance in emotional representation across datasets introduced challenges in standardizing the training process, particularly when reducing the emotion categories to 4 and then to 5 specific emotions. Additionally, the process of combining datasets to improve model resilience encountered challenges in aligning emotional labels and intensities, which could potentially influence the consistency of our findings. The average accuracy findings for both the 4-emotion and 5-emotion experiments consistently demonstrated enhanced performance when utilizing the combined dataset, as opposed to separately training SER on each individual dataset. The comparison between models trained on individual datasets versus a combined dataset approach revealed variations in performance, suggesting that not all datasets contribute equally to model accuracy and generalizability. *t*-SNE in **Figure 2** visualizes the feature embeddings extracted from the Whisper model, representing various emotional states as dis-

cerned from multiple speech datasets. It’s noteworthy how the datasets SAVEE, ESD, and JL_corpus form closely knit clusters, suggesting that their emotional representations in the feature space are similar to each other. On the other hand, the MELD, IEMOCAP, and CREMA-D datasets tend to group closer to one another, indicating a different but consistent internal representation of emotions within these datasets. The RAVDESS and TESS datasets stand out with their unique positioning in the plot, reflecting distinct emotion representation patterns extracted by the Whisper model. Moreover, for each emotion exclusive representation on each dataset reflects the inherent variations in emotions across datasets and its impact on generalization of SER.

5. Conclusion

This study proposes a novel approach to generalize SER model in leaving-one-speaker-out settings. Employing a leave-one-speaker-out methodology highlighted the model’s robustness in real-world scenarios, facing speaker variability. Extensive experiments were conducted by combining multiple SER datasets to train a Whisper-based model, followed by testing the model on a single speaker from each dataset, which yielded promising results, indicating a successful generalization across diverse datasets. These findings suggest that through careful dataset combination and targeted model training strategies, we can overcome some of the prevalent challenges in speech emotion recognition, paving the way for more universally applicable SER systems. In conclusion, our work yielded a generalized SER model adept at identifying emotions across a spectrum of datasets. This represents progress in the field, enabling the development of advanced emotion recognition systems capable of effectively handling diverse datasets. Future directions for emotion recognition in low-resource languages involve leveraging transfer learning and unsupervised learning approaches while considering cultural nuances and linguistic expertise.

6. References

- [1] A. Kulkarni, "Expressivity transfer in deep learning based text-to-speech synthesis. (synthèse vocale expressive basée sur un apprentissage profond)," in *University of Lorraine, France*, 2022. [Online]. Available: <https://theses.fr/2022LORR0122>
- [2] P. Zhou, X. P. Li, J. Li, and X. X. Jing, "Speech emotion recognition based on mixed MFCC," in *Applied Mechanics and Mechanical Engineering III*, 2013.
- [3] S. Lalitha, A. Mudupu, B. V. Nandyala, and R. Munagala, "Speech emotion recognition using DWT," in *IEEE International Conference on Computational Intelligence and Computing Research (ICCIC)*, 2015.
- [4] K. S. Rao, S. G. Koolagudi, and R. R. Vempada, "Emotion recognition from speech using global and local prosodic features," *International Journal of Speech Technology*, 2013.
- [5] K. Han, D. Yu, and I. Tashev, "Speech emotion recognition using deep neural network and extreme learning machine," *Proc. Interspeech*, 2014.
- [6] J. Lee and I. Tashev, "High-level feature representation using recurrent neural network for speech emotion recognition," in *Proc. Interspeech*, 2015.
- [7] A. Satt, S. Rozenberg, and R. Hoory, "Efficient Emotion Recognition from Speech Using Deep Learning on Spectrograms," in *Proc. Interspeech*, 2017.
- [8] 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*, 2021.
- [9] L.-W. Chen and A. Rudnicky, "Exploring wav2vec 2.0 fine-tuning for improved speech emotion recognition," in *ICASSP IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [10] J. Wagner, A. Triantafyllopoulos, H. Wierstorf, M. Schmitt, F. Burkhardt, F. Eyben, and B. W. Schuller, "Dawn of the transformer era in speech emotion recognition: Closing the valence gap," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [11] 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*, 2008.
- [12] N. Kurpukdee, S. Kasuriya, V. Chunwijitra, C. Wutiw WATCHAI, and P. Lamsrichan, "A study of support vector machines for emotional speech recognition," in *8th international conference of information and communication technology for embedded systems (IC-ICTES)*, 2017.
- [13] J. Wang, M. Xue, R. Culhane, E. Diao, J. Ding, and V. Tarokh, "Speech emotion recognition with dual-sequence lstm architecture," in *ICASSP IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [14] 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*, 2018.
- [15] M. Farooq, F. Hussain, N. K. Baloch, F. R. Raja, H. Yu, and Y. B. Zikria, "Impact of feature selection algorithm on speech emotion recognition using deep convolutional neural network," *Sensors*, 2020.
- [16] F. Dal Ri, F. Cifariello Ciardi, and N. Conci, "Speech emotion recognition and deep learning: An extensive validation using convolutional neural networks," *IEEE Access*, 2023.
- [17] 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*, 2014.
- [18] P. Jackson and S. Haq, "Surrey audio-visual expressed emotion (savee) database," *University of Surrey: Guildford, UK*, 2014.
- [19] K. Dupuis and M. K. Pichora-Fuller, "Toronto emotional speech set (tess)-younger talker_happy." Toronto: University of Toronto, Psychology Department, 2010., 2010.
- [20] M. Farooq, F. Hussain, N. K. Baloch, F. R. Raja, H. Yu, and Y. B. Zikria, "Impact of feature selection algorithm on speech emotion recognition using deep convolutional neural network," *Sensors (Basel, Switzerland)*, 2020.
- [21] L. Yue, P. Hu, S.-C. Chu, and J.-S. Pan, "Equilibrium optimizer for emotion classification from english speech signals," *IEEE Access*, 2023.
- [22] F. Haider, S. Pollak, P. Albert, and S. Luz, "Emotion recognition in low-resource settings: An evaluation of automatic feature selection methods," *Computer Speech & Language*, 2021.
- [23] T. T. L. DeJoli, Q. He, and W. Xie, "Audio, Speech and Vision Processing Lab Emotional Sound database (ASVP-ESD)." Zenodo, 2021.
- [24] A. Adigwe, N. Tits, K. E. Haddad, S. Ostadabbas, and T. Dutoit, "The emotional voices database: Towards controlling the emotion dimension in voice generation systems," *arXiv preprint arXiv:1806.09514*, 2018.
- [25] E. Parada-Cabaleiro, G. Costantini, A. Batliner, A. Baird, and B. Schuller, "Emofilm-a multilingual emotional speech corpus," in *Zenodo*, 2018.
- [26] J. James, L. Tian, and C. Watson, "An open source emotional speech corpus for human robot interaction applications," in *Proc. Interspeech*, 2018.
- [27] K. Zhou, B. Sisman, R. Liu, and H. Li, "Emotional voice conversion: Theory, databases and esd," *Speech Communication*, 2022.
- [28] J. Liu, Z. Liu, L. Wang, Y. Gao, L. Guo, and J. Dang, "Temporal attention convolutional network for speech emotion recognition with latent representation," in *Proc. Interspeech*, 2020.
- [29] S. E. Eskimez, Z. Duan, and W. Heintzelman, "Unsupervised learning approach to feature analysis for automatic speech emotion recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- [30] F. Bao, M. Neumann, and N. T. Vu, "CycleGAN-Based Emotion Style Transfer as Data Augmentation for Speech Emotion Recognition," in *Proc. Interspeech*, 2019.
- [31] B. Schuller, S. Reiter, R. Muller, M. Al-Hames, M. Lang, and G. Rigoll, "Speaker independent speech emotion recognition by ensemble classification," 2005.
- [32] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *ICML*, 2022.
- [33] H. Cao, R. Verma, and A. Nenkova, "Speaker-sensitive emotion recognition via ranking: Studies on acted and spontaneous speech," *Computer speech & language*, 2015.
- [34] S. Basu, J. Chakraborty, A. Bag, and M. Aftabuddin, "A review on emotion recognition using speech," in *International conference on inventive communication and computational technologies (ICICCT)*, 2017.
- [35] 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 *arXiv preprint arXiv:1810.02508*, 2018.
- [36] Z.-T. Liu, P. Xiao, D.-Y. Li, and M. Hao, "Speaker-independent speech emotion recognition based on cnn-blstm and multiple svms," in *Proc of ICIRA 2019*, H. Yu, J. Liu, L. Liu, Z. Ju, Y. Liu, and D. Zhou, Eds. Springer, 2019.
- [37] A. Fernández, S. Garcia, F. Herrera, and N. V. Chawla, "Smote for learning from imbalanced data: progress and challenges, marking the 15-year anniversary," *Journal of Artificial Intelligence Research*, 2018.
- [38] H. He, Y. Bai, E. A. Garcia, and S. Li, "Adasyn: Adaptive synthetic sampling approach for imbalanced learning," in *IEEE International Joint Conference on Neural Networks*, 2008.