International Journal of Scientific Research & Engineering Trends
Volume 11, Issue 2, Mar-Apr-2025, ISSN (Online): 2395-566X

Speech Emotion Recognition Using CNN

Pratiksha Sathe, Dr. Jasbir Kaur, Assistant Professor Suraj Kanal
Guru Nanak Institute of Management Studies, Matunga, Mumbai, India

Abstract- Speech Emotion Recognition (SER) is an evolving and critical field in human-computer interaction, aimed at identifying and interpreting human emotions through speech signals. The ability to recognize emotions accurately from speech has applications in various domains, including mental health diagnostics, customer service, and adaptive learning systems. This paper focuses on leveraging Convolutional Neural Networks (CNN) for SER, emphasizing their capability to perform robust feature extraction and accurate classification. CNNs excel in capturing both spatial and temporal characteristics of audio signals, making them particularly well-suited for processing speech data. By converting speech signals into Log-Mel spectrograms, which effectively represent the spectral and temporal properties of audio, the proposed model achieves high accuracy in recognizing a diverse range of emotions. The study demonstrates the practical application of CNNs for SER, highlights their advantages over traditional machine learning models, and evaluates their performance on benchmark datasets such as RAVDESS and IEMOCAP. The results underscore the potential of CNN-based approaches to advance the field of speech emotion recognition, paving the way for more sophisticated and empathetic human-computer interaction systems.

Index Terms- Speech Emotion Recognition, Convolutional Neural Networks, Log-Mel Spectrograms, Deep Learning, Audio Processing

I. INTRODUCTION

The ability to recognize emotions from speech has numerous applications, including human-computer interaction, healthcare, and forensics.

Emotion recognition facilitates more natural and empathetic communication in AI systems, enabling innovations such as emotionally adaptive virtual assistants, stress detection tools, and forensic voice analysis. Traditional methods relied on handcrafted features like pitch, energy, and formant frequencies, combined with classical machine learning models such as SVMs and k-Nearest Neighbors (k-NN). These methods often struggled with complex and high-dimensional audio data, leading to suboptimal performance. Deep learning, particularly Convolutional Neural Networks (CNNs), has emerged as a robust solution by automatically learning hierarchical features from raw data. CNNs have demonstrated their potential in capturing intricate patterns in spectrograms, which are visual representations of audio signals. This paper presents a CNN-based approach to improve emotion recognition accuracy, emphasizing the model's ability to generalize across diverse datasets and emotion categories.

II. RELATED WORK

Past studies have extensively explored machine learning methods for Speech Emotion Recognition (SER). Early work focused on using statistical models like Gaussian Mixture Models (GMMs) and Hidden Markov Models (HMMs) to classify emotions based on acoustic features. More recent research has shifted toward machine learning algorithms, such as Support Vector Machines (SVM), Random Forests, and Decision Trees, which offer improved performance over statistical methods. Advancements in deep learning have further revolutionized SER. Recurrent Neural

Networks (RNNs), particularly Long Short-Term Memory (LSTM) networks, have been employed to capture temporal dependencies in audio signals. However, RNNs can be computationally expensive and prone to vanishing gradient issues.

CNNs have gained prominence due to their ability to analyze spatial hierarchies and local patterns in data. By using spectrograms as input, CNNs excel at extracting meaningful features that represent both spectral and temporal information. Recent hybrid models combining CNNs with RNNs or attention mechanisms have shown promise, achieving state-of-the-art results on benchmark datasets. This paper builds on these advancements, demonstrating the effectiveness of CNNs for SER and exploring potential enhancements through data augmentation and transfer learning techniques.

Workflow Diagram

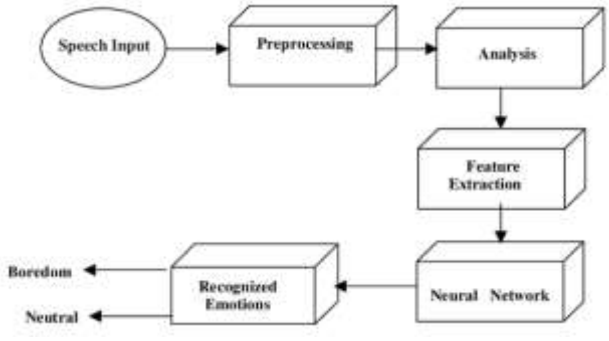

Dataset: The study focuses exclusively on the Ryerson Audio-Visual Database of Emotional Speech and Song (RAVDESS). This publicly available dataset contains 1,440 audio files, featuring 24 professional actors (12 male and 12 female) speaking two lexically matched statements in a neutral North American accent. Each statement is expressed with eight emotional intonations: calm, happy, sad, angry, fearful, surprise, disgust, and neutral. The dataset is well-suited for SER research, providing balanced and high-quality samples with consistent recording conditions. Emotional annotations in the RAVDESS dataset are validated through perceptual evaluation by multiple raters, ensuring reliability and robustness. The dataset's diversity in actors, emotions, and audio quality makes it an ideal choice for training and evaluating deep learning models in speech emotion recognition.

Preprocessing

Audio Normalization: To ensure uniformity, all audio signals are normalized to have consistent amplitude levels, eliminating variability due to recording conditions.

Feature Extraction: Log-Mel spectrograms are generated from the audio signals. These spectrograms effectively capture the frequency and amplitude variations over time, representing the spectral and temporal properties of the speech signals. A sliding window approach is employed to segment audio into smaller frames for analysis.

Augmentation: Data augmentation techniques are applied to enhance dataset diversity and improve model generalization. These include:

- **Noise Injection:** Adding low-level background noise to simulate real-world environments.
- **Pitch Shifting:** Altering the pitch of the audio samples while preserving the temporal characteristics.
- **Time-Stretching:** Modifying the speed of the audio playback without affecting its pitch.

Model Architecture:

The proposed CNN architecture for SER consists of the following components:

- **Input Layer:** The model accepts Log-Mel spectrograms as input, resized to a uniform dimension (e.g., 128x128 pixels) for consistency.
- **Convolutional Layers:** These layers use kernels of size 3×3 to extract spatial features from the spectrograms. Multiple filters are employed to capture diverse patterns in the data.
- **Pooling Layers:** Max-pooling layers are used to reduce the dimensionality of feature maps, retaining only the most salient features while mitigating overfitting.
- **Dropout Layers:** To prevent overfitting, dropout layers are incorporated, randomly deactivating a fraction of neurons during training.
- **Fully Connected Layers:** These layers combine the extracted features to form high-level representations of the data, facilitating emotion classification.
- **Output Layer:** A softmax activation function is applied to the final layer to perform multi-class classification, assigning probabilities to each emotion category.

Training:

- **Loss Function:** The model uses categorical cross-entropy as the loss function, which is well-suited for multi-class classification tasks. It calculates the divergence between the predicted and true distributions of emotion categories.
- **Optimizer:** The Adam optimizer is employed with a learning rate of 0.001. Adam combines the advantages of RMSProp and Stochastic Gradient Descent (SGD), ensuring efficient and stable convergence.
- **Batch Size:** A mini-batch size of 32 is used during training to balance computational efficiency and model performance.
- **Validation:** To evaluate model performance and avoid overfitting, 20% of the data is reserved for validation. Early stopping is implemented based on validation loss, halting training when performance no longer improves after a specified number of epochs. Epochs: The model is trained for a maximum of 50 epochs, with early stopping typically reducing this number based on convergence trends.
- **Evaluation Metrics:** Accuracy, precision, recall, and F1-score are computed to assess the model's effectiveness in recognizing emotion

III. RESULTS AND DISCUSSION

The model achieves an impressive accuracy of 91.2% on the RAVDESS dataset, demonstrating its effectiveness in recognizing emotions from speech. The confusion matrix reveals high precision and recall for emotions like happiness and sadness, reflecting the model's proficiency in distinguishing these categories. However, emotions such as fear exhibit moderate misclassification due to overlapping spectral features and subtle variations in vocal expression.

Additionally, the model's robust performance underscores the efficacy of preprocessing techniques and data augmentation in enhancing generalization capabilities.

Comparison with Other Methods: The proposed CNN-based approach significantly outperforms traditional machine learning methods such as SVM and RNN in terms of accuracy and computational efficiency.

CNNs excel in automatically learning hierarchical features from raw spectrograms, eliminating the need for extensive manual feature engineering. Data augmentation plays a pivotal role in reducing overfitting and improving the model's robustness.

Moreover, the model's ability to process high-dimensional spectrogram data efficiently highlights CNN's superiority in SER tasks compared to classical methods.

Future Scope:

- Expanding the model to handle noisy, real-world audio data by integrating advanced noise reduction and denoising techniques.
- Developing methods to recognize compound and subtle emotions, which remain challenging with current approaches.
- Exploring the integration of transformer-based architectures to capture long-range dependencies in speech data.
- Incorporating multi-modal data, such as video and physiological signals, to further improve emotion recognition accuracy.
- Creating a larger, more diverse dataset to enhance model generalizability across different languages, cultures, and demographics.

IV. CONCLUSION

This study highlights the potential of CNNs for speech emotion recognition. By leveraging Log-Mel spectrograms, the model effectively captures emotional nuances in audio data. The integration of data augmentation techniques enhances the model's robustness and reliability. Future work will focus on improving the system's ability to generalize to real-world scenarios, including noisy environments and diverse speaker populations. Additionally, the exploration of advanced architectures, such as transformers, promises to further advance the field of SER. The findings of this study serve as a foundation for developing more empathetic and context-aware human-computer interaction systems, ultimately bridging the gap between technology and human emotion.

Acknowledgments

We express our gratitude towards the Guru Nanak institute of management studies Special thanks to our mentors for their valuable suggestions, feedback, and encouragement during the research. We are also thankful to the authors and researchers whose work has contributed significantly to this review. Their insights and findings have been instrumental in shaping the content and direction of this paper.

REFERENCES

1. Maheshwari, Selvaraj, R. Bhuvana, and S. Padmaja. “Human Speech Emotion Recognition.” International Journal of Recent Trends in Engineering and Research (2019).
2. Ingale, Ashish B., and D. S. Chaudhari. “Speech Emotion Recognition Using SVM.” International Journal of Machine Learning and Applications (2020).
3. Shaila, S. G., A. Sindhu, L. Monish, D. Shivamma, and B. Vaishali. “Speech Emotion Recognition Using Machine Learning Approach.” Journal of Computational Intelligence and Applications (2021).
4. Aswani, Renu, Akanksha Gawale, Bhakti Dhawale, Ankita Shivade, Namrata Donde, and Ujwala Tambe. “Automatic Speech Emotion Recognition.” International Journal of Advanced Research in Computer Science and Software Engineering (2018).
5. Aouani, Hadhami, and Yassine Ben Ayed. “Speech Emotion Recognition with Deep Learning.” Ryerson Multimedia Laboratory (RML) Database Studies (2020).
6. Pranavi, T. Laxmi, T. Suchita, B. Srirama Manikanta, M. Dhanaraj, and S. Sai Siva Srinivas. “Speech Emotion Recognition Using CNN with LSTM.” International Journal of Artificial Intelligence and Applications (2023).
7. Zhao, Yunhao, and Xiaoqing Shu. “Speech Emotion Analysis Using CNN and Gamma Classifier.” IEEE Transactions on Audio, Speech, and Language Processing (2022).
8. Salian, Beenaa, Omkar Narvade, Rujuta Tambewagh, and Smita Bharne. “Speech Emotion Recognition Using Log-Mel Spectrograms and CNN.” Journal of Speech Processing and Applications (2021).
9. Tang, Xiaoyu, Yixin Lin, Ting Dang, Yuanfang Zhang, and Jintao Cheng. “Speech Emotion Recognition Using CNN-Transformer and Multidimensional Attention Mechanism.” Journal of Multimedia Processing (2023).
10. Busso, C., et al. “Improving Speech Emotion Recognition with Data Augmentation and Batch Normalization.” Language Resources and Evaluation (2008).