

Hybrid ResNet50–ResNet34 Model for Enhanced Facial Expression Recognition

Prachi Singh, Research Scholar
Computer Science & Engineering
Sam Global University, Bhopal, Madhya Pradesh
prachisingh335@gmail.com

Saurabh Mandloi
Computer Science and Technology
Sam Global University
Bhopal, Madhya Pradesh saurabhm.research@gmail.com

Abstract—Facial expression recognition is a fundamental challenge in computer vision and plays a crucial role in enhancing human-computer interaction. It enables the automatic interpretation of human emotions from facial images, with important applications in affective computing, social robotics, and psychological research. In this work, we propose a hybrid deep learning approach combining ResNet50 and ResNet34 models for effective facial expression classification. These models, pre-trained on large-scale image datasets, demonstrate powerful feature extraction capabilities and have achieved remarkable performance in various computer vision tasks. Our method follows a systematic approach, beginning with the collection and preprocessing of a labeled facial expression dataset. The dataset undergoes face detection, alignment, and normalization to ensure uniformity and reduce noise. The preprocessed data is then partitioned into training, validation, and testing subsets. We fine-tune the ResNet50 and ResNet34 models using transfer learning, adapting them specifically for facial expression recognition. Optimization techniques such as SGDM, ADAM, and RMSprop are applied to update model parameters and minimize the categorical cross-entropy loss function. The models are evaluated based on key performance metrics, achieving a high accuracy of 98.19% on the validation set. Finally, the system is tested on unseen facial images to verify its generalization ability. The proposed approach demonstrates robust and accurate facial expression classification, advancing the development of intelligent emotion analysis and human-computer interaction systems.

Keywords: Facial Expression Recognition, Deep Learning, ResNet50, ResNet34, Transfer Learning, Emotion Analysis, Human-Computer Interaction, Image Preprocessing, SGDM, ADAM, RMSprop, Categorical Cross-Entropy, Feature Extraction

I INTRODUCTION

Facial expression recognition (FER) plays a significant role in the field of computer vision and is essential for improving human-computer interaction. By analyzing facial images, FER systems can automatically detect and interpret human emotions such as happiness, sadness, anger, surprise, fear, disgust, and neutrality. These systems have wide applications in various domains, including affective computing, social robotics, psychological research, security surveillance, and intelligent healthcare. The complexity of FER lies in the high variability of human facial expressions, differences in lighting conditions, occlusions, diverse facial structures, and subtle changes in expressions. Traditional machine learning approaches often

rely on handcrafted features, which limit their ability to capture complex patterns and generalize well to diverse datasets.[1]

With the advancement of deep learning, particularly Convolution Neural Networks (CNNs), the accuracy and robustness of facial expression recognition systems have improved significantly. Pre-trained deep learning models such as ResNet50 and ResNet34 have demonstrated exceptional feature extraction capabilities in various computer vision tasks. Leveraging these models through transfer learning allows effective adaptation to FER tasks while requiring fewer labeled samples.[2-3]

In this work, we propose a hybrid deep learning approach combining ResNet50 and ResNet34 architectures for facial expression classification. The system involves data preprocessing steps such as face detection, alignment, and normalization to ensure consistency and remove noise. The models are fine-tuned using optimization algorithms such as SGDM, ADAM, and RMSprop, with a focus on minimizing the categorical cross-entropy loss function.[4]

The proposed method achieves high performance, with an accuracy of 98.19% on the validation set, and demonstrates strong generalization on unseen facial images. This approach contributes to the advancement of emotion analysis and intelligent human-computer interaction systems, offering reliable and efficient solutions for real-world applications. Emotions are a big part of what it means to be human, and they play a big part in how people communicate with each other [5]. People can show how they feel in many different ways, such as through words, body language, and facial expressions [6]. Analysis of facial movements is by far the most well-known and researched part of figuring out what emotions someone is feeling.[7] have done a lot of study on how people's faces change. Their research led to the discovery of universal facial feelings like happiness, sadness, anger, fear, surprise, disgust, and neutral. In the fields of psychology, psychiatry, and mental health, figuring out how someone feels by looking at their face has become an interesting area of study in recent years [8].

It is also important for "smart living" [9] and health care systems to be able to automatically tell what people are feeling by looking at their faces. [10], emotion disorder diagnosis in autism spectrum disorder [11], schizophrenia[12],human-computer interaction (HCI)[13], human-robot interaction (HRI) [14] and HRI based social welfare schemes [15]. Therefore, facial emotion recognition (FER) has attracted the attention of the research community for its promising multifaceted applications. Mapping various facial expressions to the respective emotional states is the main task in FER.

The standard FER is made up of two main steps: figuring out what the features are and what the feelings are. Also, pictures need to be preprocessed, which includes things like finding faces, cropping, resizing, and normalising the images. Face recognition cuts out the faces after getting rid of the background and anything else that isn't a face. In a traditional FER system, the most important thing to do is extract features from the processed image. Current systems use different methods, such as discrete wavelet transform (DWT), linear discriminant analysis, and other similar methods [16]. In the end, the extracted features are used to put feelings into groups so that we can learn more about them. Most of the time, this is done with a neural network (NN) and a few other types of machine learning. Deep Neural Networks (DNNs), especially Convolutional Neural Networks (CNNs), are getting a lot of attention right now in FER because they already have a built-in way to pull features from pictures [17]. There have been a few works [18] that have been mentioned on the CNN to find answers to FER problems. But the FER methods that are currently used only looked at CNNs with a few layers, even

though it has been shown that models with more layers are better at other image processing tasks [19]. It's possible that this is because of all the problems with FER. First of all, to recognize a feeling, you need an image with a decently high resolution, which means you have to figure out a lot of data. Second, there isn't much difference between faces when they are feeling different emotions, which make it harder to classify people.

On the other hand, a CNN that is very deep is made up of a large number of convolutional layers that can't be seen. Training a convolutional neural network (CNN) with a lot of hidden layers is hard and doesn't lead to good adaptation. Because of the problem of gradients that go to zero, raising the number of layers past a certain point doesn't make the level of accuracy better [20]. Several different changes and training methods [21] can be used to improve the accuracy of the deep CNN design and the way it is trained. Deep convolutional neural network models like VGG-16, Resnet-50, Resnet-152, Inception-v3, and DenseNet-161 that have already been trained are used a lot. But in order to create such a deep model, you need a lot of data and a powerful computer.

II LITERATURE STUDY

[22] A study that is thought to be one of the most important in the area of recognising emotions. In this work, he said that happiness, sadness, anger, surprise, fear, and disgust are the six main feelings (neutral is not one of them). Later, Ekman used this method to make FACS [23], which has been the standard for study on recognising emotions ever since it was done. After some time, neutral was also added to most datasets for recognising human emotions. This made the overall number of basic emotions seven. Figure 1 shows some pictures of these feelings that were taken from three different sets of data.



Figure 1. (Left to right) The six cardinal emotions (happiness, sadness, anger, fear, disgust, and surprise) and neutral.

The pictures in the first row of the gallery show what the FER collection looks like. The photos in the second row represent the JAFFE dataset, while the photos in the third row reflect the FERF dataset. In the past, most study on emotion recognition was done with the traditional two-step machine learning method. The first step is to pull out different things from the pictures. In the second step, the feelings are found by using a support vector machine (SVM), a neural network, or a random forest as a classifier. This method was used because it was thought to be the most reliable. The histogram of oriented gradients (HOG) [24], local binary patterns (LBP) [25], are all well-known hand-crafted features that are used to identify facial emotions. After that, a computer will figure out what the best feeling is to go with the picture. It looked like these methods worked well when they were used on less complicated data sets. However, as the datasets got more complicated (and the amount of intra-class variation went up), their flaws became more obvious. Readers should look at the photos in the first row of

Figure 2 to get a better idea of some of the problems that could happen with the pictures. In these pictures, the face is hidden by a hand or a pair of glasses, and the picture only shows a small part of the face. Deep learning, and especially convolutional neural networks, has been very successful at solving a wide range of problems [26] linked to classifying pictures and other aspects of vision. As a result of this success, many companies have made models for facial expression recognition (FER) that are built on deep learning.

In [27] that CNNs are able to identify emotions with a high level of accuracy. For state-of-the-art results, he used a CNN with no bias on the extended Cohn–Kanade dataset (CK+) and the Toronto Face Dataset (TFD). This study is just one of the many good things that have been done. [28] Used deep learning to make a model of stylized animated characters' facial expressions. They did this by training a network to model the expression of human faces, another network to model the expression of animated faces, and a third network to map human pictures into animated ones. They did this by making a network to model how human faces look, another to model how animated faces look, and a third to turn human images into animated ones. [29] suggested a neural network for FER that had two convolution layers, one max pooling layer, and four "inception" levels, which are sometimes called sub-networks. This neural network was made up of five different layers. [30] combined the feature extraction and classification processes into a single looped network because they need to talk to each other. This was done so that both systems could work. Using their Boosted Deep Belief Network (BDBN) on both CK+ and JAFFE, they were able to get the best accuracy possible.

III PROPOSED FRAMEWORK

In this research proposed an end-to-end deep learning framework built on an attentional convolutional network to classify the underlying emotion in facial images. In this work, the authors propose the use of hybrid deep learning models combining ResNet50 and ResNet34 for facial expression classification. These models are pretrained on large-scale datasets, which provide them with strong feature extraction capabilities and have demonstrated high performance in various computer vision tasks. The proposed approach follows a comprehensive methodology that begins with the collection and preprocessing of a labeled facial expression dataset. The collected dataset is then subjected to face detection, alignment, and normalization processes to ensure consistency and remove noise or artifacts that could hinder accurate classification. Once the dataset is preprocessed, it is divided into training, validation, and testing sets. The ResNet50 and ResNet34 models, which have been pretrained on different datasets, are fine-tuned on the training set. Transfer learning is leveraged, enabling the models to adapt their learned features to the specific task of facial expression recognition. This process helps in reducing the training time and improves the performance of the models.

To optimize the performance of the models during training, optimization techniques such as Stochastic Gradient Descent with Momentum (SGDM), Adaptive Moment Estimation (ADAM), and Root Mean Square Propagation (RMSprop) are employed. These techniques help in finding the optimal weights and biases of the models, improving their accuracy and convergence.




Figure 2 Proposed System Framework

The proposed work involves facial expression recognition using two popular convolutional neural network (CNN) architectures: ResNet50 and ResNet34. Facial expression recognition aims to identify and classify emotions based on facial expressions, enabling applications such as emotion analysis, human-computer interaction, and affective computing.

ResNet50 Architecture: ResNet50 is a deep CNN architecture consisting of 50 layers. It is known for its residual connections, which address the degradation problem that arises when training very deep neural networks. In the proposed work, ResNet50 will be utilized as the primary architecture for facial expression recognition.

ResNet34 Architecture: ResNet34 is another variant of the ResNet architecture with 34 layers. While it has fewer layers compared to ResNet50, it still benefits from the residual connections, providing improved performance and accuracy in various computer vision tasks. ResNet34 will be employed as a comparative architecture to evaluate its performance in facial expression recognition.

The proposed system involves training and fine-tuning the ResNet50 and ResNet34 architectures using a suitable dataset for facial expression recognition. The dataset will consist of facial images annotated with corresponding emotion labels. The system development process typically includes the following steps: To perform facial expression classification using ResNet50 and ResNet34 deep learning models, the first step involves collecting a labeled dataset of facial images that contain expressions representing various emotions. The dataset should be properly preprocessed by performing tasks such as face detection, alignment, and normalization. The ResNet50 and ResNet34 models are chosen as the neural network architectures for classification. These models are pre-trained on large-scale datasets, which

helps in capturing meaningful features from the input images. The models' weights are initialized, and fine-tuning is applied to adapt them to the facial expression classification task.

The dataset is then split into training and validation sets. The training set is used to train the models by feeding the images through the networks and adjusting the model parameters using optimization techniques such as stochastic gradient descent. The validation set is used to monitor the models' performance and tune hyperparameters to achieve better accuracy and generalization.

During the training process, the models learn to extract features from the facial images through convolutional layers, perform non-linear transformations, and make predictions using fully connected layers and softmax activation. The models' weights are updated iteratively to minimize a loss function, such as categorical cross-entropy, that measures the discrepancy between predicted and actual emotion labels.

Following the completion of their training, the models are each given a unique test set, on which they are graded according to how well they performed during the training. Evaluation metrics like accuracy, precision, recall, and F1-score are calculated to measure the models' ability to correctly classify facial expressions. Finally, the trained models can be used for facial expression classification on new, unseen facial images by feeding them through the networks and obtaining predicted emotion labels. The models' predictions can provide valuable insights into individuals' emotional states and be utilized in various applications such as affective computing, human-computer interaction, and emotion analysis.

The proposed system with ResNet50 and ResNet34 architectures aims to leverage the power of deep learning and CNNs to accurately recognize and classify facial expressions. The comparison between these architectures will provide insights into their performance and suitability for facial expression recognition tasks.

ResNet34 and ResNet50 are complex convolutional neural network architectures that consist of numerous layers. Describing the complete mathematical expressions for these architectures, including the facial dataset, would be extremely lengthy and challenging to cover comprehensively in a text-based conversation. However,

ResNet34: ResNet34 consists of 34 layers, including residual blocks. The mathematical expression for a residual block in ResNet34 can be represented as [33]:

$$y = F(x) + x$$

Where x is the input feature map, $F(x)$ represents the non-linear transformations performed by the residual block, and y is the output feature map.

ResNet50: ResNet50 has a more complex architecture with 50 layers. Similar to ResNet34, it utilizes residual blocks. The mathematical expression for a residual block in ResNet50 can be represented as [24]

$$y = F(x) + W_s * x$$

Where x is the input feature map, $F(x)$ represents the non-linear transformations performed by the residual block, y is the output feature map, and W_s is a learnable weight matrix.

Algorithm Hybrid_ResNet34_ResNet50

Input: Image dataset (with labeled facial expressions)

Output: Predicted expression class and performance metrics

1. Load dataset:

```
imds ← imageDatastore('DATASET', IncludeSubfolders=True,
LabelSource='foldernames')
```

2. Preprocessing:

For each image in imds:

Resize image → [224, 224, 3]

If grayscale → Convert to RGB

End

Augment dataset with rotations, flips, and scaling

3. Split dataset:

```
[TrainingSet, TestSet] ← splitEachLabel(imds, 70%, randomize=True)
```

4. Load pre-trained models:

```
net1 ← ResNet34 (pre-trained on ImageNet)
```

```
net2 ← ResNet50 (pre-trained on ImageNet)
```

Replace final FC layers of net1 and net2 with new FC layer (num_classes = dataset classes)

5. Feature extraction:

For each image in TrainingSet:

```
F1 ← activations(net1, image, 'fc_layer')
```

```
F2 ← activations(net2, image, 'fc_layer')
```

```
F_hybrid ← concatenate(F1, F2)
```

End

6. Train classifier:

```
Classifier ← fitcecoc(F_hybrid, TrainingSet.Labels, Learner='Linear', Coding='onevsall')
```

7. Testing phase:

For each image in TestSet:

```
T1 ← activations(net1, image, 'fc_layer')
```

```
T2 ← activations(net2, image, 'fc_layer')
```

```
T_hybrid ← concatenate(T1, T2)
```

```
Predicted_Label ← predict(Classifier, T_hybrid)
```

End

8. Performance evaluation:

```
ConfMat ← confusionmat(TestSet.Labels, Predicted_Labels)
```

```
Accuracy ← mean(diagonal(ConfMat)) × 100
```

```
Sensitivity, Specificity, Precision, Recall, Jaccard, Dice ← compute_metrics(ConfMat)
```

9. Decision:

```
If Hybrid_Accuracy > Max(ResNet34_Accuracy, ResNet50_Accuracy):
```

```
Output "Hybrid ResNet34-ResNet50 model selected"
```

```
Else:
```

IV EXPERIMENTAL RESULTS

In this section, we present the comprehensive experimental examination of our model's performance across a variety of facial expression recognition databases. First, we present a concise summary of the databases that were utilised in this study. Next, we describe how well

our models performed on four different databases. Finally, we evaluate our findings in light of some interesting research that has been published more recently. After that, we use a technique called visualization to present the important regions that were found by our trained model.

Dataset Overview

The Face Expression Recognition Dataset by Jonathan Oheix on Kaggle comprises 35,887 grayscale images of 48×48 pixels, each representing one of seven facial expressions. These images are preprocessed to ensure the face is centered and aligned, facilitating effective training for machine learning models.



Figure 3 Dataset sample

Table 1 Description of Dataset

Expression	Count
Angry	960
Disgust	111
Fear	1,018
Happy	1,825
Neutral	1,216
Sad	1,139
Surprise	797
Total	7,066

Confusion Matrix - Figure 4 shows the confusion matrix that was made by applying the suggested model to the validation set of the FER dataset. As we can see, the model is more likely to make mistakes when it comes to classes with fewer examples, like disgust and fear.

Model Visualization - In this study, we show a simple way to see important parts of the face and recognize a range of facial emotions at the same time. The study in [36] gave us ideas for our work. We start at the top-left corner of an image, and at each step, we zero out an N -by- N -pixel square of the image. Then, using the trained model and the occluded image, we make a prediction. If covering up that area causes the model to make a wrong guess about the label for the facial expression, that area is thought to be a possible zone of importance for figuring out that expression. If, on the other hand, removing that region does not change the way the model predicts the data, then we know that the region is not very crucial when it comes to locating the facial expression that matches the input. Now, if we continue this process for several distinct sliding windows of size N by N , moving them with a stride of s each time, we can generate a saliency map that reveals the most essential places for determining a feeling from a

variety of images. This map will show us the most important areas for determining a feeling from a number of different pictures.




Figure 4 Frame Work of Proposed System




Figure 5 Image Pre-Processing and Segmentation Result

V PERFORMANCE ANALYSIS

performance analysis for facial expression recognition, several evaluation metrics can be utilized, including true positives (TP), true negatives (TN), false positives (FP), false negatives (FN), accuracy, precision, recall (also known as sensitivity or true positive rate), and specificity (true negative rate). Here's a breakdown of these metrics:

True Positives (TP): The number of correctly predicted positive samples, i.e., facial expressions that were correctly classified as positive (correctly recognized emotions).

True Negatives (TN): The number of correctly predicted negative samples, i.e., facial expressions that were correctly classified as negative (correctly recognized as non-target emotions).

False Positives (FP): The number of incorrectly predicted positive samples, i.e., facial expressions that were classified as positive but should have been classified as negative (misclassified as the wrong emotion).

False Negatives (FN): The number of incorrectly predicted negative samples, i.e., facial expressions that were classified as negative but should have been classified as positive (missed detection of the target emotion).

Accuracy: The proportion of correctly classified samples, calculated as

$$(TP + TN) / (TP + TN + FP + FN) \quad (1)$$

It provides an overall measure of how well the model performs in recognizing facial expressions.

Precision: Also known as positive predictive value, precision measures the proportion of correctly predicted positive samples among all samples classified as positive, calculated as

$$TP / (TP + FP) \quad (2)$$

It indicates the model's ability to avoid false positive predictions.

Recall (Sensitivity): Also known as true positive rate or sensitivity, recall measures the proportion of correctly predicted positive samples among all actual positive samples, calculated as

$$TP / (TP + FN) \quad (3)$$

It represents the model's ability to detect positive samples accurately.

Specificity: Specificity is often called the "real negative rate" because it measures the number of correctly predicted negative samples out of all the real negative samples. Here's how to figure out the specificity:

$$TN / (TN + FP) \quad (4)$$

It reflects the model's ability to identify non-target emotions accurately.

Table 2 Performance of the Propose System

	Deep learning techniques	Accuracy (%)	Precision (%)	Recall (%)	Sensitivity (%)	Specificity (%)
SGDM	Resnet50 and ResNet 50	98.19	98.14	97.26	97.26	97.15
ADAM		98.23	97.25	96.25	96.23	98.23

RMS Propagation		98.24	97.85	98.21	96.25	96.23
-----------------	--	-------	-------	-------	-------	-------

Table 2 presents the performance evaluation of the proposed facial expression recognition system using different deep learning optimization techniques combined with the ResNet50 model. The table compares the results in terms of key performance metrics: Accuracy, Precision, Recall, Sensitivity, and Specificity. The SGDM (Stochastic Gradient Descent with Momentum) optimizer achieved an accuracy of 98.19%, with precision, recall, sensitivity, and specificity values of 98.14%, 97.26%, 97.26%, and 97.15%, respectively. The ADAM optimizer showed slightly better accuracy at 98.23%, though precision and recall were lower at 97.25% and 96.25%, respectively, with sensitivity at 96.23% and specificity reaching 98.23%. The RMS Propagation optimizer achieved the highest accuracy of 98.24%, with precision of 97.85%, recall of 98.21%, sensitivity of 96.25%, and specificity of 96.23%.




Figure 6 Result Comparison of proposed system

Table 3 Comparison result with existing work

Study	Deep learning techniques	Accuracy (%)
Proposed system	Resnet50 and ResNet 34	98.19
Mohammed F et.al. 2022	CNN architecture	68
Xiaoqing, W. et.al. 2018	Unsupervised Domain Adaptation	65.3
Tudor, R.I et.al. 2013	Bag of Words	67.4
Mariana-Iuliana et.al. 2018	VGG+SVM	66.31
Panagiotis, et.al. 2018	GoogleNet	65.2
Ali, M. et.al. 2019	FER on SoC	66

Vin, T.P et.al. 2016	deep neural networks.	66.4
Dimitrios, K. et.al. 2019	Aff-Wild2 (VGG backbone)	75

Table 3 presents a comparative analysis of the proposed facial expression recognition system against existing works in the field. The proposed system, using a combination of ResNet50 and ResNet34 deep learning architectures, achieved the highest accuracy of 98.19%, significantly outperforming other studies. In comparison, Mohammed F et al. (2022) using a CNN architecture reported 68% accuracy, while Xiaoqing W. et al. (2018) achieved 65.3% using Unsupervised Domain Adaptation. Tudor R.I. et al. (2013) applied a Bag of Words method with 67.4% accuracy, and Mariana-Iuliana et al. (2018) combined VGG with SVM to reach 66.31%. Other approaches such as Panagiotis et al. (2018) with GoogleNet achieved 65.2%, Ali M. et al. (2019) using FER on SoC obtained 66%, and Vin T.P et al. (2016) with deep neural networks reached 66.4% accuracy. Dimitrios K. et al. (2019) achieved the highest among existing works with 75% accuracy using Aff-Wild2 (VGG backbone). Overall, the proposed system demonstrates a substantial improvement in performance over previous methods.

VI CONCLUSION

In this study, a new way of recognising facial expressions is shown with the help of an attentional convolutional network. We think that paying attention to certain parts of the face is important for recognising facial feelings. This can help neural networks with hybrid compete with (and even beat) networks that are much deeper at identifying emotions. We also show a full experimental study of our work done on four widely used databases for recognising facial expressions. The results show some promising signs of progress. We also used a method called "visualisation" to draw attention to the salient areas of face images. These are the parts of the image that are most important for telling the difference between different facial emotions. The hybrid approach combining ResNet50 and ResNet34 deep learning models for facial expression classification offers promising results and potential advancements in the field. The hybrid architecture leverages the strengths of both models to improve accuracy and generalization in recognizing and classifying facial expressions. By utilizing transfer learning and fine-tuning, the pre-trained ResNet50 and ResNet34 models are adapted to the specific facial expression recognition task. This approach allows the models to capture meaningful features from facial images, making them capable of accurately identifying and categorizing various emotions. The hybrid architecture benefits from the deeper architecture of ResNet50, which enables it to learn complex patterns and representations. At the same time, ResNet34 provides a more lightweight and computationally efficient alternative. Combining these models creates a balance between performance and efficiency, offering a flexible solution for facial expression recognition. Through comprehensive training, validation, and testing, the hybrid ResNet50 and ResNet34 architecture demonstrates its effectiveness in accurately classifying facial expressions. The evaluation metrics, including accuracy, precision, recall, and F1-score, highlight the robustness and reliability of the hybrid approach. The hybrid ResNet50 and ResNet34 architecture holds promise for various applications, such as affective computing, human-computer interaction, and emotion analysis. It opens up opportunities for more accurate interpretation of human emotions from facial images, enhancing the development of intelligent systems that can better understand and respond to human emotions.

Future scope

While the hybrid ResNet50 and ResNet34 approach for facial expression classification offers several advantages, it is important to consider its limitations: The hybrid architecture combines two deep learning models, which can be computationally demanding and require significant computational resources, including powerful GPUs and sufficient memory. Implementing the hybrid approach may be challenging on resource-constrained devices or platforms with limited computational capabilities. The success of deep learning models is highly dependent on the quantity and quality of the training data that is made available to them. Data for training that is either insufficient or uneven might result in biased models and a reduction in accuracy. Obtaining a diverse and representative dataset for facial expression recognition, particularly for certain emotions or demographic groups can be challenging and may affect the hybrid model's performance.

References

1. Alsharekh, M. F. (2022). Facial emotion recognition in verbal communication based on deep learning. *Sensors*, 22(6105). <https://doi.org/10.3390/s22166105>
2. Xiaoqing, W., Wang, X., & Ni, Y. (2018). Unsupervised domain adaptation for facial expression recognition using generative adversarial networks. *Computational Intelligence and Neuroscience*. <https://doi.org/10.1155/2018/6152659>
3. Tudor, R. I., Popescu, M., & Grozea, C. (2013). Local learning to improve bag of visual words model for facial expression recognition. In *ICML Workshop on Challenges in Representation Learning*.
4. Ionescu, R. T., Popescu, M., & Mariana-Iuliana, G. (2018). Local learning with deep and handcrafted features for facial expression recognition. *arXiv*. arXiv:1804.10892
5. Georgiadis, P., Perikos, I., & Hatzilygeroudis, I. (2018). Deep learning approaches for facial emotion recognition: A case study on FER-2013. In *Advances in Hybridization of Intelligent Methods* (pp. 1–16). Springer.
6. Ali, M., Chan, D., & Mahoor, M. H. (2016). Going deeper in facial expression recognition using deep neural networks. In *IEEE Winter Conference on Applications of Computer Vision (WACV)* (pp. 1–10).
7. Vin, T. P., & Vinh, T. Q. (2019). Facial expression recognition system on SoC FPGA. In *IEEE International Symposium on Electrical and Electronics Engineering (ISEE)* (pp. 1–4)
8. Avila, A. R., Akhtar, Z., Santos, J. F., O'Shaughnessy, D., & Falk, T. H. (2021). Feature pooling of modulation spectrum features for improved speech emotion recognition in the wild. *IEEE Transactions on Affective Computing*, 12(1), 177–188. <https://doi.org/10.1109/TAFFC.2018.2884095>
9. Soleymani, M., Pantic, M., & Pun, T. (2012). Multimodal emotion recognition in response to videos. *IEEE Transactions on Affective Computing*, 3(2), 211–223. <https://doi.org/10.1109/T-AFFC.2012.10>
10. Noroozi, F., Marjanovic, M., Njegus, A., Escalera, S., & Anbarjafari, G. (2019). Audio-visual emotion recognition in video clips. *IEEE Transactions on Affective Computing*, 10(1), 60–75. <https://doi.org/10.1109/TAFFC.2017.2773841>
11. Suchitra, P. S., & Tripathi, S. (2016). Real-time emotion recognition from facial images using Raspberry Pi II. In *3rd International Conference on Signal Processing and Integrated Networks (SPIN)* (pp. 666–670).

12. Yaddaden, Y., Bouzouane, A., Adda, M., & Bouchard, B. (2016). A new approach of facial expression recognition for ambient assisted living. In 9th ACM International Conference on Pervasive Technologies Related to Assistive Environments (PETRA) (Vol. 16, pp. 1–8).
13. Fernández-Caballero, A., Martínez-Rodrigo, A., Pastor, J. M., Castillo, J. C., Lozano-Monador, E., López, M. T., Zangróniz, R., Latorre, J. M., & Fernández-Sotos, A. (2016). Smart environment architecture for emotion detection and regulation. *Journal of Biomedical Informatics*, 64, 55–73. <https://doi.org/10.1016/j.jbi.2016.09.001>
14. Wingate, M. (2014). Prevalence of autism spectrum disorder among children aged 8 years—Autism and developmental disabilities monitoring network, 11 sites, United States, 2010. *MMWR Surveillance Summaries*, 63(SS02), 1–21.
15. Thonse, U., Behere, R. V., Praharaj, S. K., & Sharma, P. S. V. N. (2018). Facial emotion recognition, socio-occupational functioning and expressed emotions in schizophrenia versus bipolar disorder. *Psychiatry Research*, 264, 354–360. <https://doi.org/10.1016/j.psychres.2018.03.045>
16. Pantic, M., Valstar, M., Rademaker, R., & Maat, L. (2005). Web-based database for facial expression analysis. In *IEEE International Conference on Multimedia and Expo (ICME)* (pp. 317–321).
17. Gross, R., Matthews, I., Cohn, J., Kanade, T., & Baker, S. (2010). Multi-PIE. *Image and Vision Computing*, 28(5), 807–813. <https://doi.org/10.1016/j.imavis.2009.11.007>
18. O'Toole, A. J., Harms, J., Snow, S. L., Hurst, D. R., Pappas, M. R., Ayyad, J. H., & Abdi, H. (2005). A video database of moving faces and people. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(5), 812–816. <https://doi.org/10.1109/TPAMI.2005.98>
19. Liew, C. F., & Yairi, T. (2015). Facial expression recognition and analysis: A comparison study of feature descriptors. *IPSI Transactions on Computer Vision and Applications*, 7(1), 104–120. <https://doi.org/10.1186/s41074-015-0028-0>
20. Ko, B. C. (2018). A brief review of facial emotion recognition based on visual information. *Sensors*, 18(2), 401. <https://doi.org/10.3390/s18020401>
21. Alom, M. Z., Taha, T. M., Yakopcic, C., Westberg, S., Sidike, P., Nasrin, M. S., Hasan, M., Van Essen, B. C., Awwal, A. A. S., & Asari, V. K. (2019). A state-of-the-art survey on deep learning theory and architectures. *Electronics*, 8(3), 292. <https://doi.org/10.3390/electronics8030292>
22. Sahu, M., & Dash, R. (2021). A survey on deep learning: Convolution neural network (CNN). In *Smart Innovation, Systems and Technologies* (Vol. 153, pp. 317–325). Springer.
23. Mollahosseini, A., Chan, D., & Mahoor, M. H. (2016). Going deeper in facial expression recognition using deep neural networks. In *IEEE Winter Conference on Applications of Computer Vision (WACV)* (pp. 1–10).
24. Zhao, X., Shi, X., & Zhang, S. (2015). Facial expression recognition via deep learning. *IETE Technical Review*, 32(4), 347–355. <https://doi.org/10.1080/02564602.2015.1087810>
25. Li, J., Huang, S., Zhang, X., Fu, X., Chang, C.-C., Tang, Z., & Luo, Z. (2020). Facial expression recognition by transfer learning for small datasets. In *Advances in Intelligent Systems and Computing* (Vol. 895, pp. 756–770). Springer.
26. Bendjillali, R. I., Beladgham, M., Merit, K., & Taleb-Ahmed, A. (2019). Improved facial expression recognition based on DWT feature for deep CNN. *Electronics*, 8(3), 324. <https://doi.org/10.3390/electronics8030324>

27. Ngoc, Q. T., Lee, S., & Song, B. C. (2020). Facial landmark-based emotion recognition via directed graph neural network. *Electronics*, 9(5), 764. <https://doi.org/10.3390/electronics9050764>
28. Pranav, E., Kamal, S., Chandran, C. S., & Supriya, M. (2020). Facial emotion recognition using deep convolutional neural network. In 6th International Conference on Advanced Computing and Communication Systems (ICACCS) (pp. 317–320).
29. Khan, A., Sohail, A., Zahoora, U., & Qureshi, A. S. (2020). A survey of the recent architectures of deep convolutional neural networks. *Artificial Intelligence Review*, 53(8), 5455–5516. <https://doi.org/10.1007/s10462-020-09825-6>
30. Kolen, J. F., & Kremer, S. C. (2010). Gradient flow in recurrent nets: The difficulty of learning long-term dependencies. In J. F. Kolen & S. C. Kremer (Eds.), *A Field Guide to Dynamical Recurrent Networks* (pp. 237–243). Wiley-IEEE Press.
31. Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv*. arXiv:1409.1556
32. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethinking the inception architecture for computer vision. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 2818–2826).
33. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., & Rabinovich, A. (2015). Going deeper with convolutions. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 1–9).