



IJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IV **Month of publication:** April 2026

DOI: <https://doi.org/10.22214/ijraset.2026.79879>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

Local Binary Patterns with Otsu Feature Extraction for Uneven Illumination Facial Images

A Naresh Kumar¹, Kodam Sowjanya², Bojja Veekshana³, Akkinolu Rahul Goud⁴

Department of EEE, Institute of Aeronautical Engineering, Hyderabad, India

Abstract: Depression is among the most common mental health disorders across the world. Early diagnosis plays a crucial role in timely intervention. Traditional approaches are time consuming, costly and not easily accessible in those regions that are low in resource availability. In recent days automatic depression detection is gaining attention. However, many existing frameworks struggle to handle the facial images captured under uneven illumination. This research proposes a framework that integrates Otsu thresholding for normalization of adaptive illumination and Local Binary Patterns (LBP) for robust feature extraction. A two-stage automated architecture is designed to recognize and grade the severity of depression using facial images. In the first stage, a Random Forest Classifier distinguishes between depressed and undepressed images and there by follows the second stage classifier that classifies the depressed images as mild, moderate or severe level based on the texture features extracted using local binary patterns. The model is trained and evaluated on a facial image dataset, achieving 62.7% accuracy for binary classification and 54.2% for severity classification. The framework is implemented in a Django-based web platform, supporting image upload and real-time prediction, provides a scalable and non-invasive tool for initial mental health screening.

Index Terms: depression detection, local binary patterns, random forest, facial image analysis, severity classification

I. INTRODUCTION

Mental health conditions including depression have become one of the most important issues facing society at the present day and have an impact on people of all age stages. It also resulting in significant decline in quality of life and economic productivity. As the depression is often unrecognised or unaddressed the negative impact is having on the personal and public health of individuals. If not recognised at an early stage, depression may lead to serious consequences like suicide. Traditional diagnostic procedures are based mostly on interviews and questionnaires with psychological measurement of patients which have been the means of diagnosis for many years. However, these techniques are inherently subjective, time consuming and not available in low resource settings where access to professional mental health services is limited or unavailable. Because of these limitations, there is increasing interest in developing automated techniques that can assist in the early identification of depressive systems.

In recent years, researchers have explored the use of computer vision and machine learning techniques to support automated depression detection. Facial expressions and appearance provide useful information that may reflect an individual's emotional and psychological state. Several studies have investigated depression recognition using facial image analysis and deep learning models. For example, deep neural networks have been used to capture facial appearance and temporal dynamics for estimating depression severity levels [1]. Other approaches have explored distribution-based learning methods to improve the prediction of depression scores from facial data [2]. In addition, physiological signals extracted from facial videos, such as remote photoplethysmography (rPPG), have also been investigated for depression recognition [3]. Hybrid approaches combining fuzzy logic with deep learning models have further been proposed to enhance classification performance [4]. Although these methods have shown promising results, many of them rely on computationally intensive models and may be sensitive to variations in lighting conditions present in real-world facial images. Earlier studies have also proposed frameworks that use facial expression analysis to assist in depression severity assessment [5].

In spite of the progress made in automated depression detection, several challenges still remain. One of the common issues is the variation in illumination in facial images. Changes in lighting conditions can affect the appearance of facial features and may reduce the reliability of feature extraction methods. Many existing approaches rely heavily on deep learning models that require large datasets and high computational resources. In practical scenarios, especially in low-resource environments, lightweight and robust feature extraction methods are still important. Texture-based descriptors such as Local Binary Patterns (LBP) have been widely used in image analysis because they can capture local texture information and are relatively simple to compute [6]. However, the effectiveness of such methods can still be influenced by uneven illumination present in the input images.

To address these challenges, this work proposes a framework for depression severity assessment using facial image analysis. The proposed approach combines Otsu thresholding and Local Binary Pattern (LBP) feature extraction to improve the robustness of facial texture analysis under illumination conditions. Otsu thresholding is applied during preprocessing to normalize lighting variations in facial images [7]. After preprocessing, LBP is used to extract texture features that represent local facial patterns. A two-stage classification strategy is adopted in this study. In the first stage, a Random Forest classifier is used to identify whether the input image belongs to a depressed or non-depressed category. In the second stage, the images identified as depressed are further analysed to determine the severity level, which is categorized as mild, moderate, or severe. The proposed system is implemented in a web-based platform that allows users to upload facial images and obtain predictions in real time, demonstrating the potential of the approach as a supportive tool for preliminary depression screening.

II. LITERATURE REVIEW

Recent research has explored different computational approaches for detecting depression using facial information. Facial expressions, behavioural cues, and physiological signals can reflect emotional and psychological states, which makes facial image analysis a useful area for automated depression detection. Several studies have therefore applied machine learning and deep learning techniques to analyse facial data and estimate depression levels.

One line of research focuses on deep learning methods for modelling facial appearance and facial dynamics. Yang et al. proposed a framework that uses deep neural networks to encode both static facial features and temporal changes in facial expressions for automated depression diagnosis [1]. By combining appearance and motion information from facial videos, the model aimed to capture behavioural patterns related to depressive symptoms. The results showed that facial dynamics provide important information for estimating depression severity. However, the approach relies on deep neural networks, which usually require large datasets and significant computational resources.

Another study introduced a distribution-based learning approach for predicting depression scores from facial data. De Melo et al. proposed a deep distribution learning model that represents depression severity as a probability distribution rather than a single prediction value [2]. This method attempts to better capture the uncertainty that exists in psychological assessments. The model demonstrated improved prediction performance compared with traditional regression approaches. Despite these advantages, the method still depends heavily on deep convolutional networks and large training datasets.

Researchers have also explored the use of physiological signals extracted from facial videos. In one such work, remote photoplethysmography (rPPG) signals were obtained from facial regions to analyse heart rate variations that may be related to emotional states [3]. Machine learning algorithms were then applied to these signals to detect depression. This approach introduced an alternative modality for mental health analysis. Nevertheless, extracting reliable rPPG signals can be challenging because they are sensitive to motion artifacts and lighting conditions.

Hybrid techniques combining different computational methods have also been proposed to improve detection performance. For example, a study integrating fuzzy logic with deep learning models was developed for depression detection using facial expressions [4]. In this approach, convolutional neural networks were used to extract facial features, while fuzzy logic was employed to manage uncertainty in classification. Although the hybrid framework showed improved performance compared with individual models, the system still required complex neural network architectures.

Earlier research has also attempted to design frameworks that assist in assessing depression severity through facial image analysis. Pampouchidou et al. proposed a system that analyzes facial expressions and other behavioural signals to classify different levels of depression severity [5]. The framework combined image processing techniques with machine learning algorithms. While the results demonstrated the potential of facial analysis for mental health assessment, the proposed system involved multiple data sources and complex processing steps.

Apart from deep learning approaches, traditional image analysis techniques have also been widely used for feature extraction. Local Binary Patterns (LBP) is a well-known texture descriptor that captures local intensity variations in an image [6]. The LBP operator has been successfully applied in many computer vision tasks because it is computationally simple and effective in representing texture patterns. However, the performance of texture-based descriptors can be influenced by variations in illumination present in real-world images.

Illumination normalization is therefore an important step in image preprocessing. Otsu proposed a widely used thresholding method that determines an optimal threshold value by maximizing the variance between pixel intensity classes [7]. Otsu thresholding is often applied to improve image segmentation and reduce the effects of uneven lighting conditions.

Overall, previous studies demonstrate that facial image analysis can provide useful information for automated depression detection. Many existing methods rely on deep learning models or multimodal frameworks, which may require large datasets and high computational resources. In addition, illumination variations remain a challenge for reliable feature extraction from facial images. These limitations motivate the need for lightweight approaches that can handle lighting variations while maintaining meaningful facial feature representation. The present work addresses this issue by integrating Otsu thresholding with Local Binary Pattern feature extraction for depression severity assessment.

III. DATA COLLECTION AND PREPROCESSING

A. Data Collection

For the development of the proposed depression severity assessment framework, a facial image dataset was used. The dataset was obtained for research purposes and consists of facial images categorized into two primary classes: depressed and non-depressed individuals. These categories represent facial patterns associated with depressive symptoms and those without noticeable depressive indicators. To support the training and evaluation of the classification model, the dataset was organized into training and testing sets. Both sets contain two subfolders corresponding to the depressed and non-depressed categories. The training set is used to allow the machine learning model to learn distinguishing facial characteristics associated with each class. The testing set, on the other hand, contains unseen images and is used to evaluate the ability of the trained model to generalize to new data. Separating the dataset into training and testing groups helps ensure that the evaluation results reflect the true performance of the proposed system. The images included in the dataset exhibit variations in facial expressions, illumination conditions, and image quality. Such variations are common in real-world image acquisition environments and may influence the effectiveness of feature extraction techniques. Therefore, proper preprocessing procedures are necessary to prepare the images for further analysis and to reduce the influence of unwanted variations.

B. Preprocessing

Before performing feature extraction and classification, several preprocessing operations were applied to the collected facial images. Preprocessing plays an important role in improving the quality and consistency of the input data. In the first step, facial regions were identified using a face detection method. Detecting and isolating the facial area ensures that the analysis focuses only on relevant regions of the image. After face detection, all images were resized to a fixed resolution. Standardizing the image size helps maintain uniformity across the dataset and reduces computational complexity during subsequent processing stages. This step also ensures that the extracted features have consistent dimensions for all images. One of the major challenges in facial image analysis is illumination variation. Differences in lighting conditions may alter the intensity values of pixels and affect the visibility of facial textures. Uneven illumination can therefore reduce the reliability of feature extraction techniques. To address this issue, Otsu thresholding was applied during the preprocessing stage [7]. Otsu's method automatically determines an optimal threshold value by analysing the distribution of pixel intensities in the image. This process helps separate important foreground information from background regions and improves the visibility of facial structures. Following illumination normalization, the processed images are prepared for the feature extraction stage. The Local Binary Pattern (LBP) operator is applied to capture local texture information from facial regions [6]. LBP describes the relationship between neighbouring pixels and the central pixel, allowing the system to represent facial texture patterns effectively. These texture descriptors are later used as input features for the classification model.

Through these preprocessing steps, the facial images become more consistent and suitable for further analysis. The improved image quality and normalized lighting conditions enhance the reliability of the extracted features and contribute to better classification performance in the proposed depression detection framework.




Fig. 1. Data processing pipeline

IV. SYSTEM DESIGN AND IMPLEMENTATION

The proposed system is designed to analyse facial images and estimate depression severity using texture-based feature extraction and machine learning classification. The overall workflow of the system consists of several stages, including image acquisition, preprocessing, feature extraction, and classification. Each stage plays an important role in preparing the data and generating reliable predictions.

Initially, a facial image is provided as input to the system. Since facial images may contain background information that is not relevant to the analysis, the system first performs face detection to identify and isolate the facial region. This step ensures that the subsequent processing focuses only on facial features that may reflect emotional or psychological states. After detecting the facial region, the image undergoes preprocessing to improve its quality and consistency. Images collected from different sources often exhibit variations in lighting conditions. Such variations may affect the visibility of facial textures and reduce the reliability of feature extraction. To address this issue, Otsu thresholding is applied as an illumination normalization technique [7]. This method determines an optimal threshold value based on the intensity distribution of the image, allowing the system to reduce the influence of uneven lighting.

Once the illumination variations are reduced, the system proceeds to the feature extraction stage. In this stage, Local Binary Patterns (LBP) are used to extract texture-based features from the facial image [6]. The LBP operator works by comparing the intensity value of a central pixel with its neighbouring pixels and encoding the result as a binary pattern. This process produces a numerical representation of local texture information. These texture descriptors are effective in capturing subtle variations in facial patterns that may be associated with emotional states.

The extracted LBP features are then provided as input to the classification stage. In this study, a Random Forest classifier is used because of its ability to handle high-dimensional feature spaces and provide stable classification results. Random Forest is an ensemble learning method that constructs multiple decision trees during training and combines their predictions to produce the final output. The model reduces overfitting and improves classification accuracy by aggregating the results of several decision trees.

The classification process is performed in two stages. In the first stage, the Random Forest classifier determines whether the input facial image belongs to a depressed or non-depressed category. This stage acts as a preliminary screening step. If the image is classified as non-depressed, the system directly outputs the result. However, if the image is classified as depressed, the system proceeds to the second stage of classification.

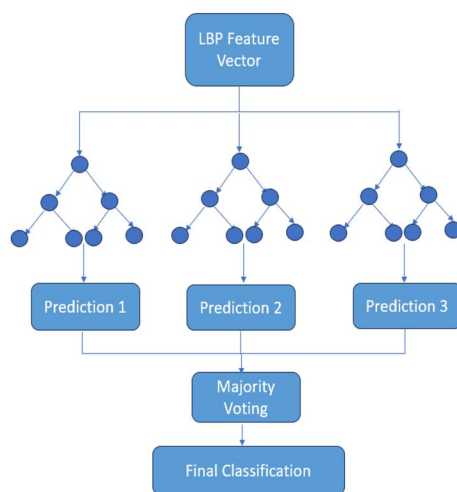


Fig. 2. Random Forest-Based Severity Classification

In the second stage, the system evaluates the severity level of depression using the extracted features. The severity is categorized into three levels: mild, moderate, and severe. By dividing the classification process into two stages, the system can first identify the presence of depressive symptoms and then estimate the level of severity more accurately. The architecture of the Random Forest classifier used in the proposed system is illustrated in Fig. 2. The figure shows how multiple decision trees are constructed from the extracted feature set and how their individual predictions are combined to produce the final classification result. Each tree in the forest independently analyses the input feature vector, and the final decision is obtained through majority voting among the trees.

Overall, the proposed system integrates preprocessing, texture-based feature extraction, and machine learning classification to provide a structured approach for depression severity assessment using facial images. The combination of Otsu thresholding and LBP feature extraction helps improve feature reliability under varying illumination conditions, while the Random Forest classifier enables effective classification of depression categories.

V. RESULTS AND DISCUSSION

The performance of the proposed depression detection framework was evaluated using the testing portion of the facial image dataset. The evaluation was carried out in two stages. In the first stage, the model identifies whether a facial image belongs to a depressed or non-depressed category. In the second stage, images identified as depressed are further classified into different severity levels. The evaluation was performed using standard classification metrics such as accuracy, precision, recall, and F1-score.

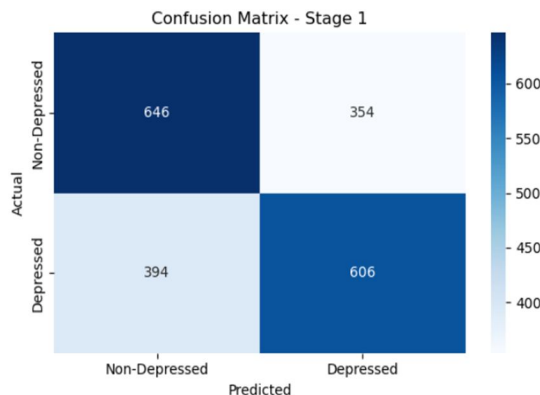


Fig. 3. Confusion Matrix at stage1

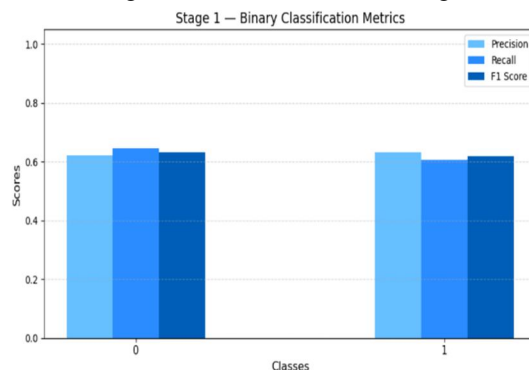


Fig. 4. Performance metrics at stage1

The first stage of the system focuses on detecting the presence of depression from facial images. A Random Forest classifier was trained using the texture features extracted through the Local Binary Pattern method. The confusion matrix obtained from the testing dataset is shown in Fig. 3. The matrix illustrates the number of correctly and incorrectly classified samples for both depressed and non-depressed classes. From the confusion matrix, it can be observed that the model is able to correctly classify a significant number of samples in both categories. However, a small number of misclassifications are also present. These errors may occur because facial expressions related to depression can sometimes appear similar to neutral expressions, especially in images captured under varying lighting conditions.

To further analyse the classification performance, several evaluation metrics were calculated. The overall performance metrics for the first stage are presented in Fig. 4. The model achieved an accuracy of approximately 63% in distinguishing between depressed and non-depressed facial images. The precision and recall values indicate that the classifier is able to identify depression-related patterns to a reasonable extent, although there is still room for improvement. The results suggest that texture-based features extracted using LBP can capture meaningful facial patterns associated with depressive behaviour.

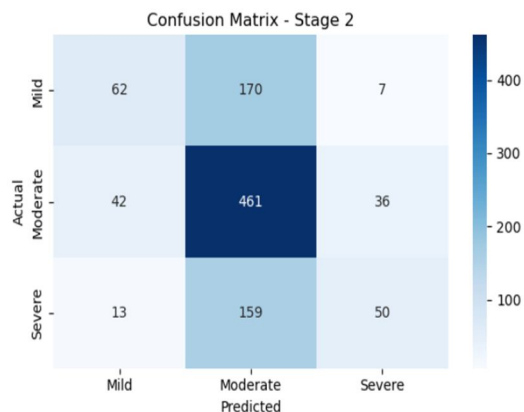


Fig. 5. Confusion Matrix at stage2




Fig. 6. Performance metrics at stage2

In the second stage of the system, the facial images identified as depressed are further analysed to determine the severity level of depression. The classification categories considered in this stage include mild, moderate, and severe depression levels. This step provides a more detailed assessment of the emotional state represented in the facial images.

The confusion matrix for the severity classification stage is shown in Fig. 5. The matrix illustrates how the classifier distributes predictions across the three severity categories. The results show that some overlap exists between the classes, which is expected because facial expressions representing different levels of depression can be subtle and difficult to distinguish. The performance metrics corresponding to the second stage classification are illustrated in Fig. 6. The model achieved an overall accuracy of approximately 57% for severity level prediction. Although the performance is lower than the binary classification stage, the results still demonstrate the capability of the proposed framework to identify different levels of depressive symptoms from facial images. The reduction in accuracy may be attributed to the increased complexity of multi-class classification and the similarity between certain facial patterns.

The experimental results demonstrate that the proposed framework is capable of identifying depression-related patterns in facial images using texture-based features and machine learning classification. The use of Otsu thresholding during preprocessing helps reduce the influence of uneven illumination, which improves the consistency of extracted features. In addition, the Local Binary Pattern descriptor provides a simple yet effective representation of local facial textures.

Compared with many existing approaches that rely heavily on deep learning architectures, the proposed system offers a relatively lightweight solution. The Random Forest classifier provides stable performance while requiring lower computational resources. However, the current results also indicate that detecting depression severity from facial images remains a challenging task. Variations in facial expressions, lighting conditions, and image quality can affect classification performance. Future improvements may include the use of larger datasets, integration of additional facial features, or the combination of multiple modalities such as facial dynamics and physiological signals. These improvements could help enhance the reliability of automated depression detection systems.

VI. CONCLUSION AND FUTURE WORK

This study presented a framework for detecting depression and assessing its severity using facial image analysis. The proposed system combines image preprocessing, texture-based feature extraction, and machine learning classification to identify patterns that may indicate depressive symptoms. Otsu thresholding was applied during the preprocessing stage to reduce the effect of uneven illumination in facial images. This step helped improve the visibility of facial features and made the images more suitable for analysis. After preprocessing, Local Binary Patterns were used to extract texture features from facial regions. These features were then provided as input to a Random Forest classifier.

The classification process was carried out in two stages. In the first stage, the system distinguished between depressed and non-depressed facial images. In the second stage, the images identified as depressed were further classified into different severity levels, including mild, moderate, and severe. Experimental results showed that the proposed approach can detect depression-related patterns from facial images with reasonable performance. The system achieved an accuracy of approximately 63% in the binary classification stage and 57% in the severity classification stage. Although the results are not perfect, they demonstrate the potential of using texture-based image features for preliminary depression assessment.

Compared with many existing approaches that rely heavily on deep learning architectures, the proposed method offers a relatively simple and computationally efficient alternative. The combination of illumination normalization and LBP feature extraction helps improve the robustness of the system when dealing with images captured under different lighting conditions. In addition, the use of a Random Forest classifier provides stable performance while maintaining moderate computational requirements. Despite these contributions, several limitations remain. Facial expressions alone may not fully represent the emotional state of an individual, and subtle differences between severity levels can make classification challenging. In some cases, lighting variations and image quality may still influence the extracted features. These factors can affect the overall classification performance of the system.

Future work can focus on improving the performance and robustness of the proposed framework. One possible direction is to use larger and more diverse datasets to allow the model to learn a wider range of facial patterns associated with depression. The integration of additional features such as facial landmarks, temporal facial dynamics, or physiological signals could also improve classification accuracy. Furthermore, combining texture-based methods with deep learning models may help capture more complex facial patterns. Such improvements could lead to more reliable automated systems that support early mental health screening and assist professionals in the assessment of depression.

REFERENCES

- [1] Y. Zhu, Y. Shang, Z. Shao, and G. Guo, "Automated depression diagnosis based on deep networks to encode facial appearance and dynamics," *IEEE Transactions on Affective Computing*, vol. 9, no. 4, pp. 578–584, Oct.–Dec. 2018.
- [2] W. C. de Melo, E. Granger, and A. Hadid, "Depression detection based on deep distribution learning," in *Proc. IEEE Int. Conf. Image Processing (ICIP)*, Taipei, Taiwan, 2019, pp. 4544–4548.
- [3] C. Á. Casado, M. L. Cañellas, and M. B. López, "Depression recognition using remote photoplethysmography from facial videos," *IEEE Transactions on Affective Computing*, vol. 14, no. 4, pp. 3305–3316, Oct.–Dec. 2023.
- [4] A. S. Rajawat, P. Bedi, S. B. Goyal, P. Bhaladhare, A. Aggarwal, and R. S. Singhal, "Fusion fuzzy logic and deep learning for depression detection using facial expressions," *Procedia Computer Science*, vol. 218, pp. 2795–2805, 2023.
- [5] A. Pampouchidou, K. Marias, M. Tsiknakis, P. Simos, F. Yang, and F. Meriaudeau, "Designing a framework for assisting depression severity assessment from facial image analysis," in *Proc. IEEE Int. Conf. Signal and Image Processing Applications (ICSIPA)*, Kuala Lumpur, Malaysia, 2015, pp. 578–583.
- [6] T. Ojala, M. Pietikäinen, and T. Mäenpää, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 7, pp. 971–987, Jul. 2002.
- [7] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62–66, Jan. 1979.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)