iJournals: International Journal of Software & Hardware Research in Engineering
ISSN:2347-4890, ijournals.in/ijshre
Volume 13 Issue 6, June 2025

Detection of Stroke-Induced Facial Paralysis Using a Convolutional Neural Network

Authors: Veda Dua[1]; Suraj Sharma[2]

Delhi Public School, R. K. Puram, New Delhi, India[1]; Mentor, On My Own Technology Mumbai, India[2]
veda.dua@gmail.com[1]; suraj.sharma@omotec.in[2]

DOI: 10.26821/IJSHRE.13.06.2025.130601

***Abstract*—**Stroke is one of the leading causes of mortality and disability. It can result in complications like facial paralysis which impair communication and emotional expression. Stroke cases are rising in India due to hypertension, diabetes, and cardiovascular diseases. There is an urgent need for affordable diagnostic tools, particularly in rural areas where healthcare access is limited. Diagnostic methods like clinical examinations and imaging systems, are either subjective, time-intensive, or expensive. This study focuses on addressing these challenges by developing a convolutional neural network (CNN) for detecting stroke-induced facial paralysis. The CNN model was trained on a dataset of 4,224 labeled images which were categorized into six classes representing different facial regions affected by paralysis along with severity. The CNN model consisted of multiple convolutional layers for extracting features. Max-pooling layers were used for reducing dimensionality, and dropout layers were used to avoid overfitting. The trained model achieved a validation accuracy of 97.04% with a validation loss of 0.0556. Its confusion matrix demonstrated the model's accuracy in classifying samples across various classes. A Streamlit-based web application was developed to allow users to upload or capture images of their face for detecting facial paralysis symptoms that indicate stroke. The app also recommends on the basis of severity of facial paralysis detected in the captured or uploaded image. The results highlight the developed model's practical utility in early stroke detection. This model reduces diagnostic costs and enables accessibility for people living in areas that are underserved. This model can empower individuals with a user-friendly solution for effective stroke management.

Keywords*—*Stroke-Induced Facial Paralysis, Convolutional Neural Network, Deep Learning in Healthcare, Stroke Diagnosis, Bell's Palsy Diagnosis, Early Stroke Detection, Cost-Effective Medical Solutions, Facial Asymmetry Detection.

I. INTRODUCTION

Stroke, which is one of the leading causes of death and lifelong disability, sometimes leaves survivors with impairments, such as facial paralysis. These impairments can restrict communication and emotional expression. Early detection is important for ensuring timely treatment. As there has been an increase in stroke cases in India due to increasing rates of hypertension, diabetes, and cardiovascular diseases, the risk of facial paralysis caused due to strokes has also increased. Stroke prevalence in India is between 84-262 per 100,000 people, with facial paralysis being one of the most probable impairments for those affected. **[1]** Stroke and its complications, which include facial paralysis, are more common among older adults, particularly in rural areas, as access to healthcare may be limited. In India, about 10% of those who are diagnosed with Bell's Palsy, a condition which leads to weakness and partial paralysis on one side of the face, die due to complications of the condition. **[2]** Present diagnostic systems of stroke induced facial paralysis are mainly based on clinical examinations by medical professionals. Such exams are mostly subjective, and time-consuming. Clinical systems utilizing traditional image processing tools cannot identify the subtlety of movement and muscle tone necessary for an accurate diagnosis. Electroneurography, the House-Brackmann Scale, Electromyography, and imaging systems like MRIs and CT Scans, are common diagnosis methods used in hospitals but they're expensive. **[3]** The House-Brackmann Scale which is relatively cheaper, but the diagnosis is very subjective as it depends on the assessor's experience. To address the limitations of these methods, This research focuses on the development of a deep learning model for identifying facial asymmetries, eye movement irregularities, and other indicators of facial paralysis. This convolutional neural network (CNN) based model is effective in processing high-dimensional data like images and extracting patterns from it. This custom CNN model, trained with a labeled dataset of images of people with facial paralysis caused due to strokes can be deployed in a web or mobile application to enhance accessibility. This solution could be helpful for people who don't have access to the traditional systems of diagnosis, and as there is no cost involved, it also allows users access to effective care early on. Hospitals could also incorporate this model into their existing systems of diagnosis.

II. Literature Review

Facial nerve paralysis (FNP) causes significant pain and functional impairment, and traditional diagnosis methods rely on subjective assessments by physicians, leading to inaccuracies. **Anping Song et al. [4]** proposed a classification system using a deep convolutional neural network (CNN) called Inception-DeepID-FNP (IDFNP), which provides a fast, accurate, and objective method for classifying FNP directly from facial images. The researchers trained the IDFNP CNN using a dataset of 1,049 clinical images categorized into seven classes of FNP severity, combining GoogleNet Inception v3 architecture with DeepID CNN and utilizing transfer learning to enhance performance. The model's performance was evaluated against neurologists' assessments, focusing on achieving high accuracy in classifying the severity of FNP. The IDFNP CNN achieved an overall classification accuracy of 97.5%, and demonstrated particularly high accuracy for severe cases of FNP, indicating its effectiveness in clinical settings. However, the classification accuracy for intermediate disease conditions was lower, and further research is needed to apply the IDFNP model to other facial diseases and improve its diagnostic capabilities for conditions like strokes, which share similar symptoms with FNP. Poor performance by models when identifying facial paralysis severity levels can be caused by an imbalance in datasets between the number of healthy and the number of FP cases. **Amira Gaber et al. [5]** combined undersampling, data augmentation, and threshold adjustment to improve such classification models. The researchers aimed to balance the dataset, enhancing the model's ability to learn from classes with less data. The researchers used fivefold cross-validation for model evaluation, employing an ensemble approach with multiple Support Vector Machines (SVMs) classifiers. Features were extracted from facial movements, and 375 records were analyzed. The classifiers were tested on metrics like accuracy, precision, sensitivity, F1-score, and specificity. The ensemble method showed better results compared to individual classifiers. The study successfully classified FP into seven severity categories, proving that an ensemble approach could be effective in improving facial paralysis prediction. The study noted limitations such as a small number of FP cases and the challenges of consistent facial expressions during assessments. Additionally, the reliance on 2D images and the high cost of 3D capture systems were highlighted as barriers to broader application. Facial paralysis severely hampers the functioning and the communication capabilities of patients necessitating standardized means to detect the disorder instead of subjective opinion from clinicians. Thus, with this aim in mind, **Gemma S. Parra-Dominguez et al. [6]** proposed a solution to facial paralysis detection from images, consisting of facial landmark extraction, calculation of facial measure, and a binary classification model. This study introduced a system consisting of 3 modules which include modules for identifying key facial points, the levels of asymmetry, and for using a multilayer perceptron classifier developed using Weka's suite to classify images. It was verified against public databases such as the Toronto NeuroFace database using a 10-fold cross-validation technique. Sensitivity, specificity, and average accuracy of the model were 98.29%, 99.54% and 94.06% respectively. Further validation across diverse datasets is necessary to assess wider applicability and real-world implementation. Facial nerve injury (FNI) remains a significant complication of parotid gland surgery (PGS) for benign tumors, despite advancements in intraoperative monitoring and preoperative assessments. **Oier Echaniz et al. [7]** addressed the gap in research on using machine learning (ML) for predicting FNI, hypothesizing that ML could improve prediction accuracy. Four ML algorithms—K-nearest neighbor (KNN), random forest (RF), boosted classification (BC), and linear discriminant classification (LDC)—were developed using clinical, radiological, histological, and cytological data. A retrospective analysis of PGS patients (2010–2019) was conducted, with recursive feature elimination and 5-fold cross-validation applied for model training and validation. Performance metrics, including sensitivity, specificity, F-score, and ROC-AUC, highlighted KNN and RF as particularly effective, with KNN achieving metrics above 0.9 for specificity and negative predictive value. While the findings suggest ML models can enhance evidence-based predictions, limitations such as retrospective data, low FNI event rates, and potential biases necessitate further validation through multicenter studies and consideration of additional confounding factors. The lack of a universal system for measuring outcomes in managing facial neuromotor disorders makes current methods for assessing facial symmetry challenging in clinical contexts. To address this, **Diego L Guarin et al. [8]** introduced Emotrics, a machine learning-based software that automates facial landmark identification and measurement calculations, designed to streamline use in medical settings. Emotrics processes frontal images where it automatically marks facial landmarks and scales measurements relative to a constant iris diameter of 11.77 mm. It also generates several metrics relevant to facial palsy. To validate the performance of the software, manual landmark adjustments were performed to enhance precision in the measurement of facial components before and after treatment. Emotrics can calculate multiple metrics in less than five seconds, which allows for objective data in the assessment of the outcome of treatment as well as the calibration of the severity of the disorder. On the other hand, it faces complications with facial asymmetry since it was trained on normal faces. Traditional stroke diagnosis methods like CT and MRI are costly and time-consuming, highlighting the need for a more efficient approach.

Ali Ahmad et al. [9] proposed an approach to diagnose stroke using facial images, with a convolutional neural network (CNN), which offers a quicker and cost-effective alternative. A total of 2,827 images of 40 stroke patients and 60 normal individuals were collected. During pre-processing, the authors applied operations like face detection, cropping, grayscale conversion, resizing to 64x64 pixels, and normalization. These processed images were then used to train the CNN model. The model was tested on 20 images, with performance evaluated through accuracy, precision, recall, and f-measure metrics. The CNN model demonstrated high effectiveness in distinguishing stroke from normal facial images, with an overall accuracy of 98%. Facial paralysis poses significant challenges, including risks to vision, difficulties in communication, and impaired emotional expression, with current assessments relying on subjective, static images that limit the understanding of facial nerve recovery. To address this, **Akshita Rao et al. [10]** proposed a video-based machine learning approach to dynamically evaluate facial function, aiming to enhance insights into nerve recovery and guide surgical planning for improved patient outcomes. This method classifies the type of facial palsy and the asymmetry of facial expressions through video data analysis, unlike static images. Using video data analysis reduced the need for manual input in the analysis of facial movements, thus producing a more accurate assessment in a shorter time. It used likelihood ratio tests, optimal transport theory, and Mahalanobis distances in the classification of facial palsy types, detection of asymmetry in facial expressions, and in the estimation of abnormal function severity against typical references. Results indicated detailed analysis of the patterns of recovery by video-based analysis, supporting clinicians with timely information to support decisions on treatment. Further research work could focus on model scalability and how to integrate it into clinical settings. Facial paralysis impacts daily life, requiring standardized medical tools for effective monitoring and treatment, ideally accessible for home use to aid patients. **Murugan Ponnusamy [11]** proposed a machine learning approach that uses K-means clustering to detect facial paralysis by analyzing facial landmarks in images. The methodology involved implementing the K-means algorithm in Python to cluster data related to facial points, enhancing detection accuracy. Evaluated through a 10-fold cross-validation process, the model demonstrated high robustness and accuracy in identifying facial paralysis, outperforming existing methods. Future research could explore the model's effectiveness across diverse populations. Traditional diagnosis relies on skilled clinicians and lacks automated, quantitative assessment, which can delay accurate FP diagnosis. To address this, **Sudhir Gogu et al. [12]** introduced an automatic facial paralysis recognition (autoFPR) approach that classifies individuals with FP versus healthy subjects using a four-stage machine learning framework: dataset creation, feature extraction, dimensionality reduction, and classification. Videos from the YouTube Facial Palsy and 300VW databases were used to create a dataset, extracting facial features like landmarks, action units, eye gaze, and head pose. Dimensionality reduction and classification were achieved using feature selection, principal component analysis, and support vector machines (SVM). The autoFPR system's performance was tested through five experiments, with metrics including accuracy, precision, recall, and F1 score, ultimately achieving a 97.3% accuracy rate. Although the study showed promising results, future research could expand the dataset to cover a broader range of FP cases and explore alternative machine learning methods to further enhance classification accuracy. Facial nerve paralysis assessment is typically subjective and varies among physicians, leading to grading inconsistencies. **Lisha Bai et al. [13]** proposed a Convolutional Neural Network (CNN) model to automatically extract and grade facial paralysis features, aiming for a more objective, efficient evaluation method. Using a hierarchical structure with convolutional layers, the model analyzes key facial movements, applying the LeakyReLU function for optimized training. Validated against models like AlexNet, VGG16, and Inception, and using the House-Brackmann scale, the CNN demonstrated substantial accuracy, precision, and F1 score improvements, with accuracy boosted by 20.3%. Limitations include a lack of public datasets for facial paralysis diagnosis and CNN overfitting issues.

Xuri Ge et al. [14] proposed an Adaptive Local-Global Relational Network (ALGRNet) to enhance facial action unit (AU) detection and improve severity estimation through adaptive learning and feature fusion. ALGRNet comprises the three elements: an Adaptive Region Learning Module for muscle region detection, Skip-BiLSTM, modeling AU relations, and a Feature Fusion & Refining Module to combine local and global features. Tested on AU benchmarks (BP4D and DISFA) and facial paralysis benchmark FPara, the new model ALGRNet reaches state-of-the-art effectiveness, achieving significant improvements in F1-scores and showcasing its robustness. While the results highlight its potential for objective assessments, further validation across diverse populations is needed. **Wasif Ali et al. [15]** proposed FP-VGGFace, a model leveraging transfer learning to fine-tune the pretrained VGGFace model for facial paralysis detection and grading. By merging the Massachusetts Eye and Ear Infirmary (MEEI) and YouTube Face Palsy (YFP) datasets, preprocessing data, and optimizing VGGFace for this task, the model achieved a remarkable 99.3% accuracy and F1-score, outperforming benchmarks like ResNet50 and VGG16. Despite its success, the reliance on private, limited datasets highlights the need for broader, more diverse data to ensure scalability and fairness in facial paralysis detection research. Unilateral peripheral facial paralysis (UPFP), the most common form of facial paralysis, causes facial asymmetry and is typically assessed using subjective scales like the House-Brackmann Grading System (HBGS). However, these scales often introduce inconsistencies due to interobserver and intraobserver variability. To address this issue, **Yuxi Liu et al. [16]** proposed an automated, objective grading method for UPFP based on HBGS, leveraging facial videos and machine learning models. The approach involved extracting facial images from videos, identifying 68 facial landmarks using dlib's model, and calculating specific features from the landmark coordinates. These features were used to train and test Random Forest (RF) and Support Vector Machine (SVM) classifiers on a dataset of 33 subjects. The method achieved a maximum accuracy of 88.9%, demonstrating its potential for objective assessments. However, the study was limited by the small dataset size, which may affect the generalizability of the results. Expanding the dataset in future research could further enhance the method's robustness and reliability. **Kieran Boochoon et al. [17]** explored the use of deep learning for assessing facial nerve palsy. It aimed to address the limitations of manual evaluations, which can be subjective and time-consuming. The methodology involved using AI to analyze facial movements and provide objective measurements of symmetry and function. Tools like the eFACE scale and Emotrics are mentioned as advancements in this area. The results highlight the potential for AI to provide real-time analysis and improve diagnostic accuracy but also point out challenges like data quality and biases in algorithms.

III. METHODOLOGY

A. Data Acquisition

A dataset containing 4,224 images of individuals with varying degrees of facial paralysis, categorized into six classes: 'mild_eye', 'mild_mouth', 'moderate_eye', 'moderate_mouth', 'severe_eye', and 'severe_mouth'. was sourced from Kaggle, These image labels represent the severity and facial regions where symptoms of facial paralysis are being observed. Labels like 'mild_eye' and 'moderate_eye' refer to varying degrees of paralysis affecting the eye, from slight to more noticeable impairment, while 'severe_eye' denotes significant paralysis where eye control may be nearly impossible. Similarly, 'mild_mouth' and 'moderate_mouth' indicate mild to moderate paralysis affecting the mouth, with 'severe_mouth' reflecting paralysis that hinders mouth control. The labeled images were used to train a deep learning model for detecting facial paralysis, which is a key symptom in identifying potential stroke risks. By analyzing facial symmetry, the trained models can classify the severity of paralysis and assist in early stroke detection.

B. Data Preprocessing

A string object containing the directory path where the images of the dataset were stored was created. The pathlib library was used to ensure proper handling of the file with different extensions. The glob function was used to list all the subdirectories within the dataset, which corresponded to the different classes of the dataset. The dataset contained bitmap image files. The total number of images in the dataset, including images present in all class subfolders was counted using glob. To ensure efficient training, batches of 16 images were created with each image resized to a consistent shape of 224*224 pixels. Resizing the images ensured uniformity in input size, which is important for feeding the data into machine learning models.

C. Dataset Splitting

The dataset was split into training and validation sets using TensorFlow's image_dataset_from_directory function. The validation_split parameter was set to 0.2, to allocate 20% of the dataset for validation, while the remaining 80% was allocated for training the model. From the total of 4,224 images in the dataset, 3,380 images were stored in the training set, and 844 images were kept for validation. This split ensured that the model is trained on a majority of the labeled images and evaluated on a separate set of images. Images from both the training and validation sets had a consistent image size of 224x224 pixels, and processing them in batches of 16 images allowed optimizing the training process for efficient and effective learning.

D. Model Architecture

The architecture of the convolutional neural network based model for detecting severity of face paralysis was created following a sequential approach. The model was designed to classify facial paralysis images into six categories based on the severity and the region of face where paralysis occurred. The first layer of the model is the Rescaling layer, which normalizes the pixel values of the images by scaling them to a range between 0 and 1, making the data suitable for an input to a neural network. Following this, the model includes several convolutional layers (Conv2D), each with increasing numbers of filters (16, 32, 64, and 128), to capture features from the images. These convolutional layers were paired with MaxPooling2D layers to reduce the spatial dimensions and highlight the most important features in the images. To avoid overfitting, Dropout layers with a 25% drop rate were applied after each convolutional block. After extracting features through the convolutional layers, the model flattens the output and passes it through a dense layer with 256 units and ReLU activation that allow it to learn complex representations of the features. The final layer, which is a dense layer that outputs the class name, has neurons equal to the number of classes. The important layers of the CNN model are shown in Fig 1.

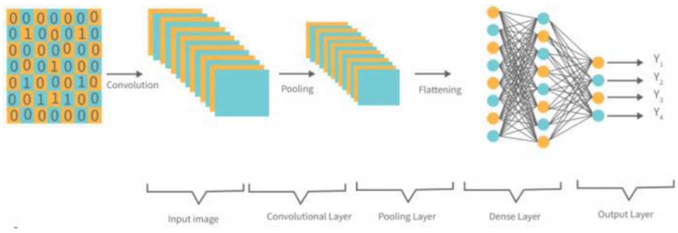

Fig 1. Layers of the CNN Model

E. Model Training

The model was compiled with the Adam optimizer, which automatically adjusts the learning rate during training to help the model learn more efficiently. The loss function used was SparseCategoricalCrossentropy, appropriate for multi-class classification tasks. The from_logits parameter was set to True to allow the raw logits to be passed directly to the loss function for classification. The evaluation metric, which was accuracy, tracked how often the model's predictions matched the true labels of the samples. The model was trained for 50 epochs, with the training data, and validation data being used to monitor its performance. The fit function returned a history object containing the loss and accuracy metrics for both the training and validation sets, which were analyzed to assess the model's progress and potential overfitting or underfitting during training.

F. Testing and Deployment

After training, the model's accuracy and loss values were plotted for both the training and validation sets across all epochs for visualizing how the model's performance improved over time, highlighting potential overfitting or underfitting trends. The accuracy plot showed consistent improvement in the model’s performance throughout the epochs, while a steady decrease in training and validation loss was observed in the loss plot. Additionally, a confusion matrix was created to assess the model's classification performance. The true labels and predicted labels from the validation set were compared, and the resulting confusion matrix provided insights into how well the model classified each of the six classes of facial paralysis dataset. A web application was also developed to deploy the trained CNN model for stroke induced facial paralysis detection using Streamlit library. The application allows users to upload or capture an image, which is then preprocessed to match the model's input dimensions (224*224 pixels). During preprocessing, the image is resized and pixel values are normalized. The model predicts the facial image's class among six categories. It also outputs confidence scores for all classes using the softmax function, and the class with the highest probability is displayed as the prediction based on which recommendations are given to the user. The app’s intuitive interface, displaying the uploaded or captured image with predictions, makes real-time stroke-related facial paralysis detection accessible.

IV. RESULTS

The trained CNN model showed excellent performance, achieving a validation accuracy of 97.04% and a validation loss of 0.0556, indicating robust generalization and classification capabilities. Graphs showing the change in training and validation metrics across all epochs revealed a consistent improvement in accuracy as shown in Fig 2, and a steady decline in loss as shown in Fig 3, demonstrating the model's effective learning without significant overfitting or underfitting. The confusion matrix which is shown in Fig. 4, demonstrates the model's strong performance in accurately categorizing images into the six facial paralysis severity classes. The predicted class displayed alongside the captured or uploaded image on the intuitive user interface developed using streamlit proved to be helpful for the users. Additionally, the suggestions or recommendations based on the predictions, further enhanced the usability of the application in diagnostic context. These results underscored the model’s practical utility and demonstrated its potential for facilitating early stroke detection through facial paralysis analysis.

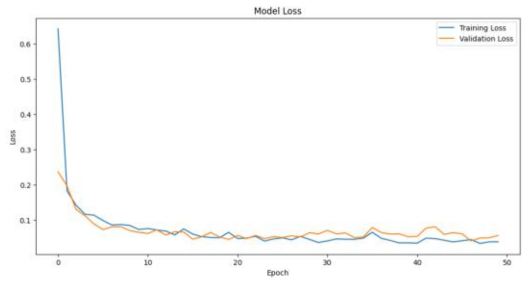

Fig 2. Decrease in Training and Validation Loss of the Facial Paralysis Detection Model Over 50 Epochs

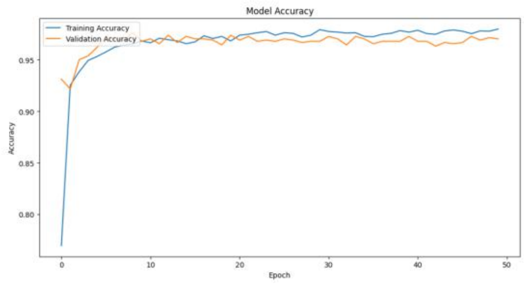

Fig 3. Improvement in Training and Validation Accuracy of the Facial Paralysis Detection Model Over 50 Epochs

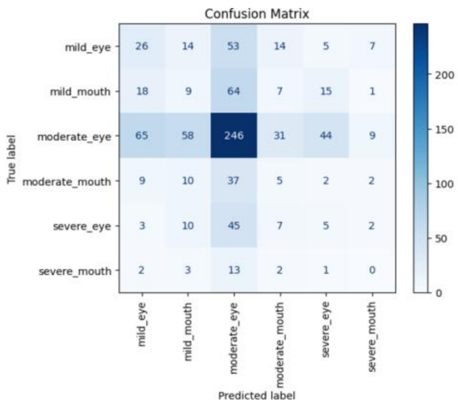

Fig 4. Confusion Matrix Depicting the Classification Accuracies for all Classes of the Dataset

VI. CONCLUSION

In conclusion, the facial paralysis detection model demonstrated exceptional performance in classifying stroke-induced facial paralysis severity. The model demonstrated effective learning and robust generalization, as evidenced by the consistent improvement in accuracy and steady decline in loss during training. The confusion matrix validated the model's ability to accurately classify images into six distinct severity classes. The deployment of the model through a Streamlit-based web application enhanced its accessibility, enabling real-time analysis of uploaded or captured images. The application's intuitive interface not only displayed predictions but also provided recommendations, proving to be a valuable tool for stroke diagnosis. The model's potential to facilitate early stroke detection through facial paralysis analysis, could be a user-friendly solution for medical professionals and individuals alike.

ACKNOWLEDGMENT

I would like to express my gratitude to the mentors Mrs. Reetu Jain and Mr. Suraj Sharma of On My Own Technology Pvt. Ltd. for extending their help in carrying out the research.

REFERENCES

[1] Pandian, Jeyaraj Durai, and Paulin Sudhan. "Stroke epidemiology and stroke care services in India." Journal of stroke 15.3 (2013): 128.

[2] Hospitals, A. (n.d.). Apollo Hospitals. Apollo Hospitals. https://www.apollohospitals.com/diseases-and-conditions/bells-palsy-symptoms-causes-risk-factors-and-treatment/#

[3] House-Brackmann Facial Paralysis Scale | Facial Paralysis Institute. (n.d.). https://facialparalysisinstitute.com/conditions/house-brackmann-grading-system/

[4] Song, Anping, et al. "Neurologist standard classification of facial nerve paralysis with deep neural networks." Future Internet 10.11 (2018): 111.

[5] Gaber, A., Taher, M.F., Wahed, M.A. et al. Classification of facial paralysis based on machine learning techniques. BioMed Eng OnLine 21, 65 (2022). https://doi.org/10.1186/s12938-022-01036-0

[6] Parra-Dominguez, Gemma S., Raul E. Sanchez-Yanez, and Carlos H. Garcia-Capulin. "Facial paralysis detection on images using key point analysis." Applied Sciences 11.5 (2021): 2435.

[7] Chiesa-Estomba, Carlos Miguel, et al. "Machine learning models for predicting facial nerve palsy in parotid gland surgery for benign tumors." journal of surgical research 262 (2021): 57-64.

[8] Guarin, Diego L., et al. "A machine learning approach for automated facial measurements in facial palsy." JAMA facial plastic surgery 20.4 (2018): 335-337.

[9] Ali, Ahmad., Muhammad, Usama., Yasir, Niaz, Khan. "Early stroke disease prediction with facial features using convolutional neural network model." IAES International Journal of Artificial Intelligence, 13 (2024).:933-933. doi: 10.11591/ijai.v13.i1.pp933-940

[10] Rao, Akshita A., Jacqueline J. Greene, and Todd P. Coleman. "Machine Learning Methods to Track Dynamic Facial Function in Facial Palsy." medRxiv (2024): 2024-03.

[11] Ponnusamy, Muruganantham. "Detection of Facial Paralysis in Human Beings Using Machine Learning Clustering." 2023 3rd International Conference on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME). IEEE, 2023.

[12] Gogu, Sridhar Reddy, and Shailesh R. Sathe. "autoFPR: An Efficient Automatic Approach for Facial Paralysis Recognition Using Facial Features." International Journal on Artificial Intelligence Tools 32.02 (2023): 2340005.

[13] Bai, Lisha, Yujian Li, and Haojie Chang. "Classification of Facial Paralysis Based on Machine Vision Techniques." Academic Journal of Computing & Information Science 7.6: 95-99.

[14] Ge, Xuri, et al. "Automatic facial paralysis estimation with facial action units." arXiv preprint arXiv:2203.01800 (2022).

[15] Ali, Wasif, et al. "A transfer learning approach for facial paralysis severity detection." IEEE Access (2023).

[16] Liu, Yuxi, et al. "Automatic assessment of facial paralysis based on facial landmarks." 2021 IEEE 2nd International Conference on Pattern Recognition and Machine Learning (PRML). IEEE, 2021.

[17] Boochoon, Kieran, et al. "Deep learning for the assessment of facial nerve palsy: opportunities and challenges." Facial Plastic Surgery 39.05 (2023): 508-511.