Oral Cancer Detection Using Deep Learning

Mrs.R. Sujatha[1], P. Venkataramana Reddy[2], Nageli Chaitanya Kumar[3], Maski Harika[4], K Kapil[5], Melingi Kavya[6]

[1,2]*Assistant Professor, Department of Artificial Intelligence and Machine Learning, Annamacharya Institute of Technology & Sciences, Tirupati, India*

[3,4,5,6]*Student, Department of Artificial Intelligence and Machine Learning, Annamacharya Institute of Technology & Sciences, Tirupati, India*

Abstract- Oral cancer is a life-threatening disease in which early identification significantly improves treatment success and patient survival. Traditional screening methods, which rely primarily on visual examination and laboratory investigations, can be time-intensive and may fail to detect early-stage abnormalities. This study presents a deep learning-based framework for automated oral cancer detection using tongue images. A transfer learning approach is employed using the DenseNet169 architecture to perform multi-class classification of oral conditions, including oral cancer, leukoplakia, thrush, lichen, hairy tongue, and healthy tongue. Data augmentation techniques were incorporated to enhance model robustness and generalization. The proposed DenseNet based model achieved an accuracy of 94.08%, significantly outperforming the conventional LeNet model. Additionally, a lightweight CNN model was developed and deployed through a Django-based web application to support real-time binary classification (Cancer / Non-Cancer). The system provides a scalable, efficient, and user-friendly solution for automated oral disease detection.

Index Terms: Oral Cancer, Deep Learning, CNN, DenseNet169, Transfer Learning, Medical Image Classification, Django, Early Detection, Artificial Intelligence in Healthcare

1. INTRODUCTION

Background and Motivation

Oral cancer is among the most prevalent cancers globally and is frequently diagnosed at advanced stages, resulting in reduced survival rates. Detecting issues early on can really boost the chances of successful treatment and better outcomes. However, conventional diagnostic procedures depend heavily on clinical expertise and laboratory-based testing.

With advancements in Artificial Intelligence, particularly Convolutional Neural Networks (CNNs), medical image analysis has become increasingly accurate and automated. Deep learning models can identify subtle patterns and visual cues that may not be easily noticeable through manual examination.

The primary motivation of this study is to enhance early diagnosis by supporting healthcare professionals with an intelligent and reliable screening system. The proposed deep learning model aims to minimize diagnostic errors by accurately identifying subtle patterns in tongue images that may be difficult to detect through manual examination.

Additionally, the system seeks to provide accessible and affordable AI-based screening tools, particularly for rural and underserved regions where advanced medical facilities and specialist expertise may be limited.

1.1 OBJECTIVES

The key objectives of this study include:

1.1.1 Development of a Deep Learning-Based Detection System

An automated classification system was developed using CNN architectures capable of analyzing tongue images.

The system operates through a structured sequence of steps to ensure accurate classification. First, the tongue image is provided as input to the model. The image then undergoes preprocessing, which includes resizing to a standard dimension and normalizing pixel values to improve training stability. After preprocessing, feature extraction is performed using convolutional layers that identify important visual patterns such as texture, color variations, and lesion boundaries. These extracted features are then passed to fully connected layers for classification. Finally, the model generates the predicted category as the output, indicating the identified oral condition.

The CNN model extracts important features such as texture variations, lesion boundaries, and color changes to differentiate between various oral conditions.

1.1.2 Implementation of Transfer Learning Using DenseNet169

Training deep neural networks from scratch requires extensive datasets, which are often limited in medical domains. Therefore, transfer learning was applied using a model pre-trained on the ImageNet dataset.

The DenseNet architecture connects each layer to every other layer, promoting efficient feature reuse and improved gradient flow. This structure reduces vanishing gradient issues and requires fewer parameters compared to traditional deep architectures.

The implementation began by loading the pre-trained DenseNet169 model without its top classification layer, allowing it to serve as a powerful feature extractor. The initial layers were then frozen to preserve the previously learned features from the ImageNet dataset. Custom dense layers were added to adapt the model specifically for oral disease classification. Finally, the last few layers were fine-tuned to improve performance on the target dataset.

This transfer learning approach significantly enhanced model efficiency and contributed to achieving an overall accuracy of 94.08%.

1.1.3 Real-Time Web-Based Deployment

To ensure practical usability, the trained model was integrated into a Django-based web application, enabling seamless real-time interaction. In this system, the user begins by uploading a tongue image through the browser interface. The uploaded image is automatically preprocessed, including resizing and normalization, before being passed to the trained model for prediction. The model then analyzes the image and generates the predicted result along with a confidence score. Finally, the outcome is displayed clearly on the interface, allowing users to receive instant screening results through a simple and accessible web platform.

BENEFITS OF THE PROPOSED SYSTEM

The developed system offers several practical advantages that enhance its real-world applicability. It provides immediate prediction results, allowing quick decision-making during screening. Since the platform is accessible through a standard web browser, it can be easily used without requiring specialized software installation.

The system is suitable for clinical environments and can assist healthcare professionals in routine examinations. Moreover, it does not require advanced technical expertise, making it user-friendly for medical staff and clinic assistants. These advantages ensure that the project extends beyond theoretical research and functions as an effective and practical clinical support tool.

MODEL PERFORMANCE COMPARISON

Evaluating multiple architectures is essential to justify the selection of DenseNet for this application. The LeNet architecture is a basic convolutional neural network with a limited number of layers, making it suitable for simple image classification tasks. However, due to its shallow structure, it is insufficient for capturing the complex and subtle patterns present in medical images. In contrast.

DenseNet is a deep neural network with dense connectivity between layers, which promotes efficient feature reuse and improved gradient flow. This architecture enables the extraction of detailed and fine-grained patterns, making it highly effective for complex medical image analysis such as oral cancer detection.

Quantitative Performance Comparison

Metric	DenseNet	LeNet
Accuracy	94.08%	64.02%
Precision	94.16%	64.06%
Recall	94.70%	64.03%
F1-Score	94.07%	63.01%

Analysis

The analysis indicates that DenseNet effectively captures subtle disease-related features that are critical for accurate oral cancer detection. In contrast, LeNet struggles to model the complex variations present in oral lesions due to its simpler architecture. Furthermore, DenseNet demonstrates stronger generalization capability, enabling it to perform consistently across diverse image conditions. This comparative evaluation scientifically supports the selection of DenseNet as the preferred architecture for oral cancer detection.

1.2 SCOPE OF THE PROJECT

The project addresses the following major components:

1.2.1 Multi-Class Classification

The system performs multi-class classification of tongue images by categorizing them into six distinct classes: oral cancer, leukoplakia, thrush, lichen, hairy tongue, and healthy tongue.

By analyzing visual features such as texture, color variations, and lesion patterns, the model accurately distinguishes between these different oral conditions, enabling comprehensive detection and diagnosis support.

1.2.2 Binary Classification for Deployment

For efficient real-time implementation, the model is simplified into a binary classification system consisting of two categories: Cancer and Non-Cancer. By reducing the number of output classes, the system achieves faster processing and quicker predictions, which is essential for live screening applications.

This simplification also enhances reliability and stability during real-time use, making the model more practical for clinical environments and routine screening scenarios.

1.2.3 Web-Based Implementation

A web platform was developed using Django to allow users to upload images and receive instant predictions. The application performs preprocessing, classification, and result display in real time. Built-in authentication mechanisms and database support ensure secure storage of patient records and efficient user management.

1.2.4 Future Expansion Opportunities

The system can be further extended to enhance its scalability and practical impact. Future improvements may include mobile application integration, allowing users to perform screenings directly through smartphones for greater accessibility.

Cloud-based deployment can also be implemented to enable remote access, centralized data storage, and improved computational efficiency. Additionally, integrating the system with hospital information systems would facilitate seamless data sharing, patient record management, and streamlined clinical workflows, making the solution more comprehensive and adaptable to real-world healthcare environments.

II.LITERATURE SURVEY

Traditional methods of oral cancer detection primarily rely on clinical and laboratory-based procedures. Conventional diagnosis typically involves visual inspection by medical professionals to identify visible abnormalities in the oral cavity.

If suspicious lesions are observed, tissue biopsy procedures are performed, followed by microscopic histopathological examination to confirm the presence of cancerous cells. These methods are considered the standard approach for accurate diagnosis in clinical practice.

However, these traditional diagnostic techniques have several limitations. The overall process is often lengthy and time-consuming, which may delay early treatment. Additionally, biopsy and laboratory investigations can be expensive, increasing the financial burden on patients. These methods also depend heavily on experienced medical professionals and specialized laboratory infrastructure, which may not be readily available in rural or underserved regions, thereby limiting accessibility to timely diagnosis.

Advances in Deep Learning for Medical Imaging

Recent developments in deep learning field have significantly improved medical image analysis. Convolutional Neural Networks automatically learn discriminative image features without manual intervention, resulting in higher classification accuracy compared to traditional machine learning approaches.

DenseNet enhances learning efficiency through dense layer connections, enabling improved gradient flow and better feature propagation. Transfer learning further strengthens performance by utilizing knowledge from pre-trained models, reducing data requirements and computational cost.

DenseNet vs. LeNet

Compared to LeNet, DenseNet extracts deeper and more meaningful representations due to its interconnected layer structure. This architecture facilitates improved information flow and allows the model to identify subtle disease-related characteristics.

DenseNet handles variations in illumination, scale, and background conditions more effectively, making it better suited for complex medical image classification tasks.

Challenges in Oral Cancer Detection

Automated oral cancer detection faces several real-world challenges that can affect system performance and reliability. One major difficulty is the visual similarity between different oral diseases, as many conditions share overlapping features such as color changes and lesion patterns, making accurate classification more complex. Additionally, imbalanced dataset distribution can lead to biased model predictions, especially when certain disease categories have fewer samples. The limited availability of properly annotated medical images further restricts effective training of deep learning models.

Moreover, deep model training requires high computational resources, which can be costly and time-consuming. Addressing these challenges is essential to ensure accurate, robust, and reliable clinical deployment of the system.

III. METHODOLOGY

3.1 DATASET PREPARATION

The dataset used in this study consists of tongue images categorized into six different classes, including oral cancer, leukoplakia, thrush, lichen, hairy tongue, and healthy tongue. These images were collected from clinical and medical sources to ensure authenticity and relevance. Proper labeling of each image was performed to maintain classification accuracy during training. The dataset was organized into structured folders based on class categories. This structured dataset preparation ensures that the deep learning model can effectively learn disease-specific patterns.

Preprocessing Techniques

Before training the model, all images underwent preprocessing to ensure uniformity and improved model performance. Each image was resized to 180×180 pixels to maintain consistency in input dimensions for the CNN architecture. Pixel values were normalized to a range between 0 and 1 to stabilize and accelerate the training process. Noise and irrelevant variations were minimized through standard preprocessing steps.

These preprocessing techniques help improve model convergence and overall accuracy.

Data Augmentation

To enhance the robustness of the model and prevent overfitting, data augmentation techniques were applied to the training dataset. Images were randomly rotated to allow the model to learn from different viewing angles. Horizontal flipping was performed to increase dataset diversity without collecting new images. Zooming operations were also applied to help the model recognize features at different scales. These augmentation techniques effectively increased the dataset size and improved generalization performance.

Data Splitting

The prepared dataset was divided into training and validation sets to evaluate model performance effectively. Approximately 80% of the data was used for training the model, allowing it to learn underlying patterns and features. The remaining 20% was allocated for validation to monitor performance during training. This split ensures unbiased evaluation and helps detect overfitting. Proper data splitting improves the reliability and credibility of the experimental results.

3.2 SYSTEM ARCHITECTURE

The system architecture for Oral Cancer Detection Using Deep Learning consists of multiple stages, including image input, preprocessing, model training, prediction, and output display. The system begins with the collection of tongue images representing various oral conditions, such as oral cancer, leukoplakia, thrush, lichen, hairy tongue, and healthy tongue. These images serve as the input to the deep learning framework. In the preprocessing stage, all images are resized to 180×180 pixels to maintain uniformity and ensure compatibility with the model input requirements. Pixel values are scaled to help the model train more effectively. Data augmentation techniques such as horizontal flipping, rotation, and zooming are applied to increase dataset diversity and reduce overfitting.

The core of the system is the deep learning model based on DenseNet169 architecture using transfer learning. The pre-trained DenseNet model (trained on ImageNet) is fine-tuned by adding custom dense layers for classification. The model uses convolutional layers for feature extraction, pooling layers for dimensionality reduction, and fully connected layers for classification. A softmax function is used to classify the output into multiple classes. The dataset is split into 80% training and 20% validation data. The model is trained using the Adam optimizer with categorical binary cross-entropy loss.

The system's performance is analyzed using key evaluation metrics: accuracy, precision, recall, and F1-score.

In the prediction phase, when a user uploads a tongue image through the web interface, the image undergoes the same preprocessing steps before being passed to the trained model. The model predicts the class label along with a confidence score. Finally, in the output stage, the predicted result (Cancer or Non-Cancer / specific oral condition) is displayed on the user interface. The system can be integrated into a Django-based web application for real-time oral cancer detection, making it practical for clinical support systems.

3.3 DEEP LEARNING MODEL

The CNN model serves as the core of the Oral Cancer Detection system. It is responsible for automatically extracting important visual features from tongue images and classifying them into different oral conditions, such as oral cancer, leukoplakia, thrush, lichen, hairy tongue, and healthy tongue.

2.3.1 Model Architecture

A Convolutional Neural Network (CNN) was implemented using the DenseNet169 architecture with transfer learning for efficient medical image classification. The model begins with an input layer that accepts images resized to 180×180×3. The pre-trained DenseNet169 model, originally trained on the ImageNet dataset, is used as the base for feature extraction.

The convolutional layers automatically extract low-level and high-level features such as texture patterns, color variations, and lesion boundaries from tongue images. Dense connections in DenseNet improve feature reuse and reduce the vanishing gradient problem. After feature extraction, a Global Average Pooling layer is applied to reduce dimensionality. The output is then passed through fully connected (dense) layers with ReLU activation.

Dropout layers are added to prevent overfitting and improve generalization. The final classification layer uses a Softmax activation function for multi-class classification of oral conditions. For binary classification (Cancer / Non-Cancer), a sigmoid activation function can also be used.

2.3.2 Compilation and Training

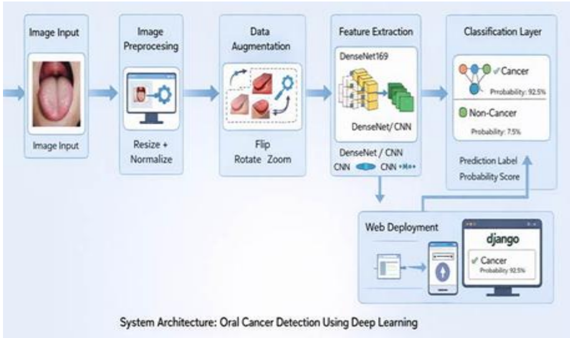

System Architecture: Oral Cancer Detection Using Deep Learning

The model was compiled using the categorical cross-entropy loss function for multi-class classification and optimized using the Adam optimizer with a learning rate of 0.001.

Accuracy was used as the primary evaluation metric to measure model performance. The training process was conducted for 50 epochs with a batch size of 32 to ensure stable convergence. Transfer learning significantly reduced training time and improved performance compared to training a CNN from scratch. During training, validation accuracy and loss were continuously monitored to ensure optimal learning and avoid overfitting.

3.4 TRAINING AND VALIDATION

The dataset was divided into 80% training data and 20% validation data to ensure unbiased performance evaluation. Data augmentation techniques such as horizontal flipping, rotation, and zooming were applied during training to increase dataset diversity and improve model robustness.

The validation set was used to monitor model accuracy and loss after each epoch. Early stopping and learning rate reduction callbacks were implemented to enhance model performance. The final model achieved high accuracy on both training and validation sets. Performance across different oral disease classes was evaluated using a confusion matrix and classification report, including precision, recall, and F1-score metrics.

3.5 USER INTERFACE

The user interface of the Oral Cancer Detection system is designed to be simple, interactive, and user-friendly, especially for non-technical users such as healthcare staff and clinic assistants. The interface allows users to upload a tongue image for oral cancer analysis. It provides clear buttons such as “Upload Image” and “Predict” to guide users through the process. The interface is clean and responsive, ensuring smooth performance across laptops, desktops, and mobile devices. It automatically adjusts to different screen sizes for better usability. Proper error-handling mechanisms are implemented to manage unsupported image formats or poor-quality images by displaying appropriate feedback messages.

Once the prediction is completed, the system displays the predicted class (Cancer / Non-Cancer or specific oral condition) along with the confidence score, making the results easy to understand.

IV. IMPLEMENTATION

4.1 TOOLS AND TECHNOLOGIES

To develop the Oral Cancer Detection system using Deep Learning, a combination of modern tools and technologies was used to ensure accuracy, performance, and scalability. Python was used as the primary programming language due to its simplicity and extensive support for machine learning libraries. TensorFlow served as the deep learning framework for model development and training, while Keras was used as a high-level API for designing and fine-tuning the DenseNet169 architecture. Image preprocessing tasks such as resizing, normalization, and augmentation were performed using OpenCV and Pillow (PIL). NumPy was used for efficient handling of image arrays and mathematical operations. Matplotlib was used to visualize training performance, including accuracy and loss curves. For deployment, a Django-based web application was developed to enable real-time prediction. The system stores and manages data using SQLite as the backend database. This integrated technology stack ensures efficient training, testing, and deployment of the oral cancer detection model.

4 .2 CODE OVERVIEW

The implementation of the Oral Cancer Detection system using Deep Learning is divided into three main modules:

4.2.1 Loading Data and Preprocessing

The dataset is loaded using TensorFlow and processed using OpenCV and Pillow. All tongue images are converted to RGB format, resized to 180×180 pixels, and normalized by scaling pixel values between 0 and 1. Data augmentation techniques such as rotation, flipping, and zooming are applied to increase dataset diversity. The dataset is then converted into NumPy arrays, and labels are encoded for classification. Finally, the dataset is split into 80% training data and 20% validation data to ensure proper model evaluation.

4.2.2 Constructing and Training the CNN Model

The DenseNet169 architecture is used as the base model with transfer learning. The pre-trained model is loaded without the top classification layer. Custom dense layers with ReLU activation and dropout layers are added to improve classification performance and reduce overfitting. The model is compiled using categorical cross-entropy loss (for multi-class classification) and optimized using the Adam optimizer. Model training is performed for 50 epochs utilizing a batch size of 32 to optimize learning efficiency.

After training, the model is saved as a .h5 file for future predictions and deployment.

4.2.3 Prediction and Classification

In the prediction phase, when a user uploads a tongue image through the Django web interface, the image undergoes the same preprocessing steps (resizing and normalization). The processed image is passed to the trained model for classification. The model predicts the probability of each class and displays the final prediction along with the confidence score on the web interface. This allows real-time oral cancer detection and makes the system suitable for clinical support applications.

V. RESULT AND DISCUSSION

5.1MODEL PERFORMANCE

During the testing phase, the proposed DenseNet169-based deep learning model achieved a validation accuracy of 94.08% after training for 50 epochs. The training and validation accuracy curves showed steady improvement across epochs, indicating effective learning without significant overfitting. Similarly, the training and validation loss curves gradually decreased, demonstrating stable convergence of the model. Transfer learning played a major role in improving performance, as the pre-trained DenseNet169 model was fine-tuned on the oral cancer dataset. Data augmentation techniques such as rotation, horizontal flipping, and zooming were applied during training to improve generalization and robustness. These techniques helped the model adapt to variations in tongue orientation, lighting conditions, and image quality.

The model was evaluated using standard performance metrics, which includes Accuracy, Precision, F1-Score, and Recall. The DenseNet169 model achieved:

· Accuracy : 94.08%

· Precision : 94.16%

· Recall : 94.7%

· F1-Score : 94.07%

To compare model performance, a LeNet architecture was also implemented. The LeNet model achieved significantly lower performance:

· Accuracy : 64.02%

· Precision : 64.06%

· Recall : 64.03%

· F1-Score : 63.01%

This comparison clearly demonstrates that the DenseNet architecture is more suitable for complex medical image classification tasks such as oral cancer detection. Dense connections allow better feature reuse and deeper feature extraction, enabling the model to identify subtle lesion patterns in tongue images. A confusion matrix was generated to evaluate class-wise performance. The results showed high true positive rates for oral cancer and healthy classes, with minimal misclassification between similar oral conditions such as leukoplakia and lichen.

Minor errors occurred in visually similar disease categories, which can be further reduced with a larger dataset and enhanced augmentation strategies. Overall, the experimental results confirm that the proposed deep learning-based system provides reliable and accurate detection of oral cancer and can be effectively used as a clinical decision-support tool.

Figure 2: Training and Validation Accuracy

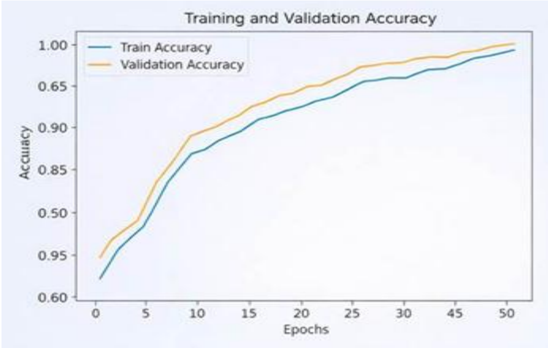

Figure 3: Confusion Matrix

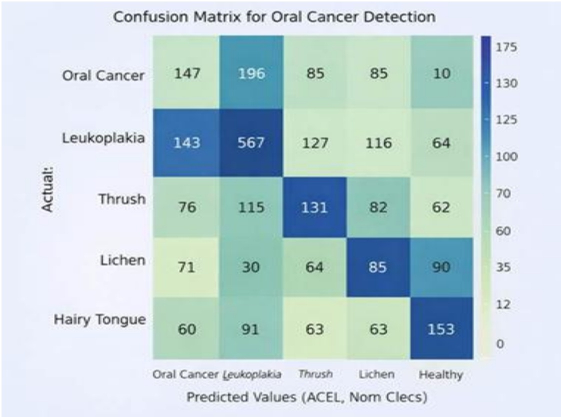

5.2 Output Images:

During the testing phase, the Oral Cancer Detection system demonstrated strong generalization capability, achieving a validation accuracy of 94.08% after 50 epochs of training. The training and validation accuracy curves showed stable convergence, indicating that the model learned meaningful patterns without significant overfitting. Data augmentation techniques such as rotation, horizontal flipping, and zooming improved robustness by allowing the model to handle variations in tongue orientation, lighting conditions, and image quality. During the testing phase, the Oral Cancer Detection system demonstrated strong generalization capability, achieving a validation accuracy of 94.08% after 50 epochs of training. The training and validation accuracy curves showed stable convergence, indicating that the model learned meaningful patterns without significant overfitting. Data augmentation techniques such as rotation, horizontal flipping, and zooming improved robustness by allowing the model to handle variations in tongue orientation, lighting conditions, and image quality.

The confusion matrix analysis revealed high precision and recall across most oral disease classes. Minor misclassifications occurred between visually similar conditions, such as leukoplakia and lichen, due to overlapping visual features.

Despite this, the system maintains a high reliability level, providing consistent and accurate predictions suitable for clinical decision-support applications.

5.3 COMPARISON WITH TRADITIONAL METHODS

Traditional oral cancer detection methods rely heavily on manual clinical examination, biopsy procedures, and histopathological analysis. These methods are time-consuming, costly, and require expert medical professionals. Additionally, early-stage oral cancer lesions may be visually subtle and difficult to identify through manual inspection alone.

In contrast, the proposed deep learning-based system uses Convolutional Neural Networks (CNNs) to automatically extract hierarchical features directly from tongue images. The DenseNet169 architecture enhances feature reuse and improves gradient flow, allowing the model to capture complex lesion patterns more effectively than traditional image processing techniques.

Unlike traditional handcrafted feature extraction methods, CNN-based approaches eliminate manual intervention and improve scalability. The deep learning model demonstrates higher accuracy, better generalization, and faster processing time, making it more suitable for automated screening and early detection. This shift toward AI-driven diagnosis significantly enhances the efficiency and reliability of oral cancer detection systems.

5.4 FUTURE WORK

Collecting more varied and larger datasets can strengthen the model's performance and reliability. A more balanced dataset will also help address class imbalance and reduce bias during training. Integration of Explainable AI (XAI) techniques, such as Grad-CAM, can help visualize important regions in tongue images, improving interpretability and trust among healthcare professionals.

The proposed system may be further enhanced by integrating real-time camera-based screening capabilities within mobile and desktop applications. Deployment on cloud platforms would facilitate remote access, particularly benefiting rural healthcare centers and using ensemble learning methods and combining patient history with image data can further improve diagnostic accuracy.

Moreover, large-scale clinical validation using real-world datasets and collaboration with healthcare institutions would strengthen the reliability, robustness, and practical applicability of the system in real medical settings.

VI. CONCLUSION

The study validates the effectiveness of deep learning techniques in enabling early detection of oral cancer through tongue image analysis. By implementing the DenseNet169 architecture with transfer learning, the system achieved high classification performance, attaining an accuracy of 94.08%, along with strong precision, recall, and F1-score metrics. The results show that advanced CNN models can accurately detect small visual differences between oral cancer and similar oral conditions

The incorporation of data augmentation and systematic preprocessing techniques significantly enhanced model generalization and robustness against variations in lighting conditions, orientation, and image quality. When compared to traditional diagnostic approaches and simpler CNN models such as LeNet, the proposed DenseNet-based framework demonstrated superior feature extraction capability and substantially improved classification performance.

Furthermore, integrating the trained model into a Django-based web application enhances its practical applicability by enabling real-time screening. The user-friendly interface supports seamless image upload and instant prediction, making it suitable for clinical assistance and preliminary diagnostic evaluation.

Overall, the proposed Oral Cancer Detection system provides an efficient and reliable method for automated oral disease classification. With further clinical validation and larger real-world datasets, it has strong potential to assist healthcare professionals in early diagnosis and improve patient survival rates.

REFERENCES

[1] A. Ramola, A. K. Shakya and A. Bergman, “Comprehensive analysis of advancement in optical biosensing techniques for early detection of cancerous cells,” Biosensors, vol. 15, no. 5, p. 292, 2025. DOI: https://doi.org/10.3390/bios15050292

[2] T. Patten and P. Jacobs, “Natural-language processing,” IEEE Expert, vol. 9, no. 1, pp. 35–39, Feb. 1994. DOI: https://doi.org/10.1109/64.294351

[3] S. W. Chang, S. Abdul-Kareem, A. F. Merican and R. B. Zain, “Oral cancer prognosis based on clinicopathologic and genomic markers using a hybrid of feature selection and machine learning methods,” BMC Bioinformatics, vol. 14, Art. no. 170, May 2013. DOI: https://doi.org/10.1186/1471-2105-14-170

[4] D. I. Conway, M. Purkayastha and I. G. Chestnutt, “The changing epidemiology of oral cancer: Definitions, trends, and risk factors,” British Dental Journal, vol. 225, no. 9, pp. 867–873, Nov. 2018. DOI: https://doi.org/10.1038/sj.bdj.2018.1024

[5] N. Issa, J. Leonas, B. C. Jham, J. C. Mitchell and M. C. Cuevas-Nunez, “Early detection of oral potentially malignant disorders using machine learning: A retrospective pilot study,” General Dentistry, vol. 70, no. 6, pp. 60–64, Nov.–Dec. 2022. DOI: https://doi.org/10.18326/gdj.2022.70.6.60

[6] E. Albalawi et al., “Oral squamous cell carcinoma detection using EfficientNet on histopathological images,” Frontiers in Medicine, vol. 10, 2024. DOI: https://doi.org/10.3389/fmed.2023.1349336

[7] A. R. Omran, “The epidemiologic transition: A theory of the epidemiology of population change,” Milbank Memorial Fund Quarterly, 1971. DOI: https://doi.org/10.2307/3349375

[8] S. Warnakulasuriya, “Global epidemiology of oral and oropharyngeal cancer,” Oral Oncology, vol. 45, no. 4–5, pp. 309–316, 2009. DOI: https://doi.org/10.1016/j.oraloncology.2008.06.002

[9] K. Vayadande et al., “Innovative approaches for skin disease identification in machine learning: A comprehensive study,” Oral Oncology Reports, vol. 10, 2024. DOI: https://doi.org/10.1016/j.oor.2024.100365

[10] G. Vennila et al., “Enhancing speech emotion recognition using deep learning networks on live calls,” Proceedings of the 5th International Conference for Emerging Technology (INCET), Karnataka, India, May 2024. DOI: https://doi.org/10.1109/INCET61531.2024.10593786

[11] V. Sanh, L. Debut, J. Chaumond and T. Wolf, “DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter,” arXiv preprint, 2019. DOI: https://doi.org/10.48550/arXiv.1910.01108

[12] R. O. Alabi, A. Almangush, M. Elmusrati and A. A. Mäkitie, “Deep machine learning for oral cancer: From precise diagnosis to precision

medicine," Frontiers in Oral Health, vol. 2, 2022. DOI: https://doi.org/10.3389/froh.2021.794248
[13] D. P. Harsha and Hemanth Kumar, "Counting individuals in an image using machine learning technique," International Journal of Innovative Science and Research Technology, 2024. DOI: https://doi.org/10.38124/ijisrt/IJISRT24JUL1266
[14] R. Zebari and A. Sallow, "Face detection and recognition using OpenCV," Journal of Soft Computing and Data Mining, vol. 2, 2021. DOI: https://doi.org/10.30880/jscdm.2021.02.02.008
[15] M. W. Lingen, J. R. Kalmar, T. Karrison and P. M. Speight, "Critical evaluation of diagnostic aids for the detection of oral cancer," Oral Oncology, vol. 44, no. 1, pp. 10–22, 2008. DOI: https://doi.org/10.1016/j.oraloncology.2007.06.011
[16] K. Dhanuthai et al., "Oral cancer: A multicenter study," Medicina Oral Patologia Oral y Cirugia Bucal, vol. 23, no. 1, pp. e23–e29, 2018. DOI: https://doi.org/10.4317/medoral.21999
[17] S. Petti, "Lifestyle risk factors for oral cancer," Oral Oncology, vol. 45, no. 4–5, pp. 340–350, 2009. DOI: https://doi.org/10.1016/j.oraloncology.2008.05.018
[18] M. Hashibe et al., "Alcohol drinking in never users of tobacco and risk of head and neck cancer," Journal of the National Cancer Institute, vol. 99, pp. 777–789. DOI: https://doi.org/10.1093/jnci/djk179
[19] J. Gigliotti, S. Madathil and N. Makhoul, "Delays in oral cavity cancer," International Journal of Oral and Maxillofacial Surgery, vol. 48, pp. 1131–1137, 2019. DOI: https://doi.org/10.1016/j.ijom.2019.02.012
[20] R. Chavva, S. Jeba Priya and Dr. Mathu, "Oral Cancer Detection Using Deep Learning," Proceedings of the International Conference on Science Technology Engineering and Management (ICSTEM), 2024. DOI: https://doi.org/10.1109/ICSTEM61137.2024.10561172