

# Deep Learning Based Oral Disease Detection System

**G. Sangeetha lakshmi<sup>1</sup>, P. Dhanalakshmi<sup>2</sup>**

*G. Sangeetha lakshmi<sup>1</sup>, M.S(IT), M.Phil., Assistant Professor and head,*

*Department of Computer Science and Applications*

*D.K.M. College for Women (Autonomous)*

*P. Dhanalakshmi<sup>2</sup>, PG Student, Department of Computer Science and Applications*

*Department of Computer Science and Applications*

*D.K.M. College for Women (Autonomous), Vellore, Tamil Nadu, India*

## Abstract

Oral diseases represent a significant global health concern and may lead to severe complications if not detected at an early stage. Common conditions such as dental caries, gingivitis, oral ulcers, calculus, hypodontia, and tooth discoloration are largely preventable and treatable when timely diagnosis is ensured. However, traditional diagnostic methods rely heavily on physical clinical examinations performed by dental professionals, which are often inaccessible in rural and underserved areas. This paper presents a Deep Learning Based Oral Disease Detection System — a web-based intelligent application that employs a Convolutional Neural Network (CNN) integrated with transfer learning (EfficientNetB0) for automated multi-class oral disease classification. Oral cavity images uploaded by users undergo preprocessing operations including resizing to 224×224 pixels, normalization, and data augmentation prior to analysis. The model was trained on a dataset of 6,200 images across seven disease categories and achieved an overall classification accuracy of 84% on the test set, with class-specific F1-scores reaching up to 0.99 for hypodontia and tooth discoloration. The system is deployed via the Streamlit framework and provides real-time disease predictions along with disease-specific preventive and treatment guidelines. The proposed system serves as an accessible screening tool that bridges the gap between limited dental healthcare access and the need for early diagnosis, particularly in remote and resource-constrained regions.

**Keywords:** *Oral Disease Detection, Deep Learning, Convolutional Neural Network, EfficientNetB0, Transfer Learning, Medical Image Classification, Streamlit.*

## I. INTRODUCTION

Oral diseases are among the most prevalent health conditions affecting people of all age groups worldwide. According to the World Health Organization (WHO), oral diseases affect nearly 3.5 billion people globally, making them one of the most widespread non-communicable disease burdens. Conditions such as dental caries, gingivitis, mouth ulcers, calculus, hypodontia, and tooth discoloration, if left undetected and untreated, can progress to severe systemic complications including cardiovascular disease and compromised nutritional intake. Early detection plays a crucial role in improving treatment outcomes, reducing long-term medical expenditure, and increasing patient survival rates.

Traditional oral disease diagnosis depends heavily on physical clinical examinations by trained dental professionals, which require specialized facilities that are often inaccessible in rural and underserved communities. The shortage of dental practitioners in low- and middle-income countries, combined with the high cost of specialist

consultations, means that a significant proportion of oral conditions go undiagnosed until they reach advanced stages. This diagnostic gap underscores the urgent need for automated, accessible, and cost-effective screening tools that can support early identification of oral disease.

Recent advances in Artificial Intelligence (AI), particularly Deep Learning, have demonstrated significant potential in medical image analysis. Convolutional Neural Networks (CNNs) are capable of automatically learning discriminative features from images and performing accurate multi-class classification without handcrafted feature engineering. Transfer learning, which involves fine-tuning models pretrained on large datasets, has further accelerated the adoption of deep learning in medical imaging by enabling high accuracy even with relatively small disease-specific datasets.

This paper presents a Deep Learning Based Oral Disease Detection System implemented as a web-based application using the Streamlit framework. The system enables users to upload oral cavity images and receive real-time disease predictions through a trained CNN model built on the EfficientNetB0 architecture. The system classifies images into seven categories and generates disease-specific preventive measures and treatment guidelines. The proposed system aims to support early screening and improve access to preliminary diagnosis, particularly for individuals in underserved regions where trained dental professionals are scarce.

## II. RELATED WORK

The application of deep learning to oral and dental disease detection has gained considerable momentum in recent years. Several studies have explored CNN-based approaches for detecting a variety of oral conditions, with varying model architectures, datasets, and deployment strategies.

Reddy and Narayanan (2021) introduced an AI-powered oral disease identification system targeting remote healthcare environments. Using deep learning on mobile-captured oral images, the system successfully detected mouth ulcers, tooth discoloration, and gum infections, emphasizing the role of AI in enabling telemedicine-based early diagnosis.

Park and Lee (2022) proposed a multi-class CNN classification framework for detecting caries, calculus, gingivitis, hypodontia, and healthy teeth. Their results confirmed the suitability of deep learning for real-world

multi-label oral disease screening with strong performance across all evaluated categories.

Sharma and Gupta (2022) explored fine-tuning of pretrained CNN architectures such as VGG16 and ResNet50 on oral image datasets, demonstrating that transfer learning significantly reduces training time while maintaining high accuracy, even with limited labelled data.

Ahmed and Rahman (2023) presented an end-to-end system combining preprocessing, automated feature extraction, and CNN classification into a single pipeline, accurately identifying dental caries, mouth ulcers, and gingivitis. Mehta and Joshi (2024) extended this by coupling AI inference with a web interface for browser-based image submission and instant prediction. Iyer and Mohan (2025) further demonstrated robustness across diverse patient demographics using ensemble deep learning strategies.

A systematic review by Sivari et al. (2023) surveyed deep learning methods for dental anomaly and disease diagnosis, concluding that CNN-based architectures with transfer learning consistently outperform traditional machine learning on oral image classification benchmarks. The review identified gingivitis and early-stage caries as the most challenging categories due to their visual similarity with other conditions — a finding corroborated by the results of the present study.

### III. METHODOLOGY

#### A. Dataset Collection and Distribution

The dataset used in this study comprises 6,200 oral cavity images collected from publicly available dental health datasets, online oral health repositories, and peer-reviewed research publications. The images span seven disease categories: Calculus, Caries, Data Caries, Gingivitis, Mouth Ulcer, Tooth Discoloration, and Hypodontia. Images were captured under varying lighting conditions using different camera devices, representing patients from diverse demographic backgrounds to ensure the model generalizes across real-world imaging variations. The dataset was partitioned into training (70%), validation (20%), and testing (10%) subsets, as summarized in Table III.

**TABLE III**  
*Dataset Class Distribution*

Disease Class	Train	Val.	Test
Calculus	805	230	115
Caries	882	252	126
Data Caries	364	104	52
Gingivitis	448	128	64
Mouth Ulcer	483	138	69
Tooth Discolor.	448	128	64
Hypodontia	525	150	75
<b>Total</b>	<b>3955</b>	<b>1130</b>	<b>565</b>

#### B. Image Preprocessing

All input images were resized to 224×224 pixels to conform to the input requirements of the EfficientNetB0 backbone. Pixel values were normalized to the [0, 1] range by dividing

each channel value by 255, ensuring stable gradient flow and faster convergence during training. Data augmentation techniques — including random rotation (up to 20°), horizontal flipping, zoom (up to 15%), and brightness adjustment — were applied exclusively to the training set to increase sample diversity and improve model robustness. In select cases, background regions were masked to direct the network's attention toward clinically relevant oral structures, reducing noise in the feature extraction stage.

#### C. Model Architecture and Transfer Learning

The proposed system employs EfficientNetB0 as the feature extraction backbone, selected for its compound scaling approach that balances network depth, width, and resolution, yielding superior accuracy per parameter compared to other architectures of similar complexity. EfficientNetB0 was initialized with weights pretrained on ImageNet and subsequently fine-tuned on the oral disease dataset using transfer learning. The convolutional layers extract hierarchical feature representations capturing low-level texture gradients, mid-level structural patterns, and high-level semantic features corresponding to disease-specific tissue characteristics such as lesion boundaries, surface deposits, and gingival inflammation.

The 1,280-dimensional feature vectors from EfficientNetB0's global average pooling layer are passed through a fully connected dense layer (256 neurons, ReLU activation), a dropout layer (rate = 0.4) for regularization, and a Softmax output layer with seven neurons. The model was compiled using Adam optimizer (learning rate = 0.001) with categorical cross-entropy loss. Training ran for five epochs with early stopping (patience = 3) and learning rate reduction on plateau (factor = 0.5) callbacks.

#### D. System Architecture

The overall system pipeline integrates five modules: (1) Oral Image Input Module — users submit images via the Streamlit web interface; (2) Image Preprocessing Module — resizing, normalization, and format conversion are applied; (3) Feature Extraction Module — EfficientNetB0 convolutional layers generate compact feature vectors; (4) Classification Module — fully connected layers and Softmax produce a probability distribution across the seven disease classes; and (5) Result Display Module — the predicted disease label, confidence score, and health recommendations are rendered to the user.

#### E. Disease Recommendation Module

Following successful classification, the system retrieves disease-specific recommendations from a structured dictionary mapped to each of the seven classes. Recommendations are organized into Treatment (e.g., professional scaling for calculus; topical analgesic ointments for mouth ulcers; dental implants or braces for hypodontia) and Prevention (e.g., reducing dietary sugar for caries; avoiding smoking for gingivitis; limiting tea and coffee for tooth discoloration). This module empowers users with actionable health information and encourages timely professional consultation.

### IV. IMPLEMENTATION

The system is implemented entirely in Python using TensorFlow 2.13 and Keras for model development and inference, OpenCV and PIL for image handling, and Streamlit for the web application layer. The trained

EfficientNetB0 model is serialized in HDF5 format and loaded at application startup for low-latency inference.

The application incorporates a JSON-based user authentication system with separate registration and login pages. Upon successful login, users access the detection interface, upload an oral image, and receive predictions in real time. The uploaded image is decoded, converted to BGR format, resized to 150×150 pixels, reshaped into a batch tensor of shape (1, 150, 150, 3), and passed to the model's predict method. The argmax of the probability output determines the predicted class label, which retrieves the corresponding recommendations.

Supporting libraries include NumPy for numerical array operations, Pandas for structured data handling, Seaborn and Matplotlib for training curve visualization, and Scikit-learn for generating the classification report and confusion matrix. The application runs locally with 'streamlit run main.py' and is compatible with standard consumer-grade hardware without requiring GPU acceleration for inference, making it deployable in resource-constrained environments.

## V. RESULTS AND DISCUSSION

The proposed model was evaluated on a held-out test set of 565 images spanning all seven disease categories. Performance was assessed using precision, recall, F1-score, and overall accuracy. Table I presents per-class metrics, Table II summarizes aggregate performance, and the confusion matrix was analysed to understand the distribution of misclassifications.

**TABLE I**  
*Per-Class Classification Performance*

Disease Class	Prec.	Recall	F1	Support
Calculus	0.82	0.85	0.84	115
Caries	0.80	0.71	0.75	126
Data Caries	0.98	0.98	0.98	52
Gingivitis	0.42	0.48	0.45	64
Mouth Ulcer	0.95	1.00	0.97	69
Tooth Discolor.	1.00	0.97	0.98	64
Hypodontia	0.99	0.99	0.99	75

**TABLE II**  
*Overall Model Performance Summary*

Metric	Value
Overall Accuracy	84%
Macro Avg Precision	0.85
Macro Avg Recall	0.85
Macro Avg F1-Score	0.85
Test Set Size	565 images
Training Epochs	5
Backbone	EfficientNetB0

The model achieved an overall accuracy of 84% across the seven-class oral disease classification task. Hypodontia achieved the highest performance (F1 = 0.99), followed by tooth discoloration (F1 = 0.98), data caries (F1 = 0.98), and mouth ulcer (F1 = 0.97). Calculus (F1 = 0.84) and caries (F1 = 0.75) performed moderately well, while gingivitis exhibited the lowest F1-score of 0.45, attributable to its visual similarity with early-stage calculus and caries in clinical oral images — a well-documented challenge in oral disease image analysis.

The training and validation accuracy curves demonstrated consistent improvement across the five training epochs. Training accuracy progressed from approximately 70% at epoch 1 to 90% by epoch 5, while validation accuracy stabilized around 85%, indicating adequate generalization without significant overfitting. The corresponding training and validation loss curves showed steady monotonic convergence, confirming the effectiveness of the EfficientNetB0 backbone with transfer learning and the appropriateness of the Adam optimizer configuration.

Analysis of the confusion matrix revealed strong diagonal dominance across all seven classes, confirming that the majority of test samples were correctly classified. The primary source of misclassification was inter-class confusion between gingivitis, calculus, and caries — conditions that share overlapping visual characteristics including gingival redness, surface deposit accumulation, and tissue inflammation. Of the 64 gingivitis test samples, 13 were misclassified as calculus and 19 as caries. Improving discrimination for these visually ambiguous classes through class-specific data augmentation, attention-based mechanisms, or multi-scale feature fusion represents a key direction for future work.

A comparison with related methods is provided in Table IV. The proposed EfficientNetB0-based system achieves competitive accuracy relative to existing approaches while additionally offering a fully integrated, web-deployable interface with disease-specific recommendation generation — features not present in most prior work.

**TABLE IV**  
*Comparison with Existing Methods*

Method	Architecture	Classes	Accuracy
Park & Lee (2022)	Custom CNN	5	~82%
Sharma & Gupta (2022)	VGG16 / ResNet	4	~86%
Ahmed & Rahman (2023)	CNN Pipeline	3	~80%
Mehta & Joshi (2024)	Web-CNN	5	~83%
<b>Proposed System</b>	<b>EfficientNetB0</b>	<b>7</b>	<b>84%</b>

Inference time per image was measured at under 2 seconds on standard CPU hardware (Intel Core i5, 8GB RAM), enabling near-real-time prediction suitable for practical deployment without requiring specialized GPU infrastructure. This confirms that the system meets the

performance requirements of a clinical screening tool intended for low-resource environments.

## VI. SYSTEM INTERFACE AND APPLICATIONS

The system is accessible through a Streamlit web application featuring four primary screens. The Registration Page allows new users to create accounts with a username and password, with credentials stored persistently in a local JSON file. The Login Page authenticates existing users and redirects them to the main detection interface. The Image Upload Page provides authenticated users with a drag-and-drop or file-browser interface supporting JPEG, PNG, and BMP formats, with a file size limit of 200MB. The Result Page displays the uploaded image, the predicted disease condition with confidence score, and tailored treatment and preventive health recommendations corresponding to the identified disease class.

The application is designed to be deployable across a range of use cases: in dental clinics as a preliminary screening aid; within telemedicine platforms for remote oral health assessment; in community health awareness programs; and as a self-assessment tool for individuals in areas with limited dental care access. The interface requires no technical expertise from the end user, making it accessible to patients, community health workers, and non-specialist healthcare providers alike. Future enhancements to the interface may include multilingual support, confidence threshold warnings, and integration with electronic health record (EHR) systems.

## VII. CONCLUSION

This paper presented a Deep Learning Based Oral Disease Detection System that leverages EfficientNetB0 with transfer learning for automated multi-class classification of oral diseases from oral cavity images. Trained on a dataset of 6,200 images across seven disease categories, the system achieved an overall test accuracy of 84%, with several classes — including hypodontia, tooth discoloration, data caries, and mouth ulcer — attaining F1-scores of 0.97 or above. The web-based implementation via Streamlit delivers a user-friendly, real-time diagnostic interface that does not require GPU infrastructure, making it practically deployable in resource-constrained and rural settings where trained dental professionals are scarce.

The system effectively addresses the critical gap between limited dental healthcare access and the need for timely oral disease screening, contributing meaningfully to early detection and preventive healthcare. The integrated disease-specific recommendation module further enhances its value by empowering users with actionable treatment and prevention guidance.

Future work will focus on several directions: expanding the dataset to include greater demographic and imaging device diversity; addressing gingivitis classification through attention mechanisms and class-specific augmentation; exploring advanced deep learning architectures such as ResNet50V2, DenseNet, and Vision Transformers (ViT) for improved accuracy; incorporating multi-modal data including patient history and clinical symptoms for richer diagnostic support; and integrating the system with telemedicine and hospital management platforms to facilitate

broader clinical adoption and longitudinal patient monitoring.

## REFERENCES

- [1] E. Sivari, G.B. Senirkentli, E. Bostanci et al., "Deep Learning in Diagnosis of Dental Anomalies and Diseases: A Systematic Review," *Diagnostics*, vol. 13, no. 15, 2023.
- [2] T. Xue, L. Chen, and Q. Sun, "Deep learning method to automatically diagnose periodontal bone loss and periodontitis stage in dental panoramic radiograph," *Journal of Dentistry*, 2024.
- [3] R. Tan, X. Zhu, S. Chen et al., "Caries lesions diagnosis with deep convolutional neural network in intraoral QLF images by handheld device," *BMC Oral Health*, vol. 24, p. 754, 2024.
- [4] Z. Yesil, I.S. Bayrakdar, E. Bilgir et al., "Evaluation of tooth development stages with deep learning-based artificial intelligence algorithm," *BMC Oral Health*, vol. 24, p. 1034, 2024.
- [5] J. Rashid, B.S. Qaisar, M. Faheem et al., "Mouth and oral disease classification using InceptionResNetV2," *Multimedia Tools and Applications*, 2024.
- [6] SmartOralDx Research Team, "SmartOralDx: A deep learning-powered system for precise classification of oral diseases from clinical imagery," *International Biomedical Journal*, 2025.
- [7] J. Park and S. Lee, "Multi-Class Oral Disease Detection Using Deep Learning Models," 2022.
- [8] N. Sharma and R. Gupta, "Computer-Aided Diagnosis of Oral Diseases Using CNN and Transfer Learning," 2022.
- [9] T. Ahmed and F. Rahman, "Intelligent Oral Disease Detection System Using Deep Learning," 2023.
- [10] P. Mehta and A. Joshi, "Web-Based Oral Disease Detection Using Deep Learning Techniques," 2024.