Oral Cancer Detection Using AI

Assist Prof.Kavita Reddy[1], Swagath H[2], Tejaswini R R[3], Shreyash Gowda B A[4], Prajwal L Nadig[5]
Department of Computer Science-Artificial Intelligence and Machine Learning, AMC engineering college, Bengaluru, India

***Abstract*—Oral cancer is one of the top causes of cancer-related deaths worldwide. This is mainly due to late diagnosis and limited access to specialized healthcare. Early detection can greatly improve survival rates and lower treatment costs. Recent progress in artificial intelligence (AI), especially in deep learning and image processing, has changed how we diagnose cancer. This study brings together findings from various reviews and experimental works that examine the role of AI in detecting oral cancer. AI algorithms, particularly convolutional neural networks(CNNs),have shown impressive diagnostic accuracy, ranging from 81% to 99.7%. They have high sensitivity and specificity when distinguishing between normal, precancerous, and malignant lesions. AI systems that use photographic, cytological, and autofluorescence imaging allow for non-invasive and real-time screening. Some of these systems even work through smartphone apps, making them accessible in low-resource areas. Combining machine learning with telemedicine has shown good agreement with expert clinical evaluations, enabling remote and early diagnosis.**

***Index Terms*—Oral Cancer, Artificial Intelligence, Deep Learning, Convolutional Neural Network (CNN), Image Processing, Early Diagnosis, Computer Vision**

I. INTRODUCTION

One of the main causes of cancer-related deaths globally is still oral cancer, which frequently goes undetected until it has progressed, reducing survival rates. Improving prognosis and treatment results depends heavily on early detection. Visual inspection and biopsy are examples of traditional diagnostic techniques that are invasive, time-consuming, and rely on professional interpretation. Our project suggests an AI-based system for early oral cancer detection in order to get around these restrictions. The system uses deep learning and image processing to accurately and non-invasively identify potentially malignant lesions in oral images. The model's use of Convolutional Neural Networks (CNNs) improves diagnostic accuracy, lowers human error, and permits economical screening—particularly in places with limited resources or remote locations.

Using deep learning models to analyze images of the oral cavity and automatically identify possible cancerous lesions is the goal of the proposed project, "Oral Cancer Detection Using AI." To improve image clarity, extract pertinent features, and classify images with high precision, the system makes use of AI algorithms and image processing techniques. By using smartphone-based imaging solutions, this method not only cuts down on diagnostic time and human error but also makes screening available in remote and underdeveloped areas.

This project aims to enable early, accurate, and non-invasive detection of oral cancer by combining artificial intelligence with oral healthcare. This will improve patient survival rates and lessen the burden of the disease worldwide. In the end, AI-driven diagnostics could revolutionize conventional medical procedures by bridging the gap between clinical medicine and cutting-edge technology to enable quicker and more accurate cancer detection.

II. RELATED WORKS

Recent developments in image processing and artificial intelligence (AI) have demonstrated impressive promise in enhancing the early diagnosis and detection of oral cancer, a condition that remains a significant global health concern because of its high mortality rates and late-stage diagnosis. Numerous studies have looked into AI's capacity to detect oral cancers, contrasting its results with those of more conventional diagnostic techniques like visual inspection and biopsy.

The usefulness of AI methods for identifying oral cancer was thoroughly examined in a review by Al-

Rawi et al. (2022) that was published in the International Dental Journal. They examined 17 separate studies with 7,245 patients and more than 69,000 photos. At accuracies ranging from 81% to 99.7%, sensitivity up to 98.75%, and specificity as high as 100%, it was discovered that deep learning (DL) models—in particular, convolutional neural networks (CNNs)—performed better than conventional machine learning (ML) techniques. The authors stressed that deep learning models are capable of processing sizable, intricate datasets and spotting minute morphological details in medical photos that the human eye might miss.

A review of AI's potential for early and non-invasive oral cancer detection was published in the journal Cancers by García-Pola et al. in 2021. They looked at 36 studies that were released between 2000 and 2020 and discovered that AI has been successfully applied to telemedicine, cytology, autofluorescence imaging, and oral photos. AI systems that were mobile and cloud-based performed well, particularly in places with a shortage of medical professionals. Combining deep learning with smartphone imaging allowed for the detection of cancerous lesions with over 85% sensitivity and up to 93% specificity. The study came to the conclusion that while standardizing data and lowering model errors are still necessary, AI can enhance large-scale detection and early screening.

A smartphone-based AI diagnostic framework that uses deep learning for oral lesion classification was more recently proposed by Mira et al. (2024). Their strategy centered on using handheld smartphone cameras to capture images in a practical and affordable manner. The authors increased diagnostic consistency by implementing a "center positioning" rule and resampling techniques to lessen image capture variability. The system achieved 84.3% accuracy, 83% sensitivity, and 96.6% specificity using a CNN-based model that was trained on 455 test images that represented five different oral conditions. The growing viability of AI-powered mobile diagnostic tools for point-of-care screening, particularly in underserved or rural areas, is highlighted by this study.

III. IMPLEMENTATION

The developed CNN based oral cancer detection framework comprises five core modules:

- Data preprocessing and augmentation
- Deep feature extraction using CNN
- Model Training and optimization
- Image classification
- Performance evaluation matrix

Data preprocessing and augmentation

Preprocessing of data constitutes a crucial step in developing an AI-based oral cancer detection system, which helps in the quality and consistency of input images before feeding into the neural network. It involves several operations during the preprocessing stage, like noise removal, resizing, normalization, and enhancing contrast. Oral cavity images inherently show variability in lighting conditions, orientation, and colour because of differences in acquisition sources; therefore, images from smartphone cameras, autofluorescence, or clinical imaging systems first undergo standardization to the same size, such as 224×224 pixels, and also need colour channel normalization, typically RGB. To increase the visibility of the lesions and mucosal textures, histogram equalization and adaptive contrast enhancement methods are also applied. In various studies, resampling and centering have been proposed as ways to reduce variability in images to ensure that the lesions maintain consistent positioning, enhancing model robustness and the accuracy of classification.

Data augmentation further expands the dataset for the purpose of overcoming class imbalance and preventing overfitting. The transformation includes horizontal and vertical flipping, rotation, scaling, translation, and random cropping, which can simulate real variations in capturing images. Other approaches involve brightness and contrast adjustment, Gaussian noise injection, and zoom augmentation to create more diversity within the training dataset and improve model generalization on unseen data. In the case of oral histopathological images, patch extraction is one of the techniques often used to increase the sample count while preserving details of the lesions. These augmentation strategies allow deep learning models, especially CNNs, to reach higher sensitivity and specificity in identifying normal, precancerous, and malignant tissues. Overall, careful preprocessing and augmentation enhance the reliability, robustness, and diagnostic precision of the AI-based oral cancer detection framework.

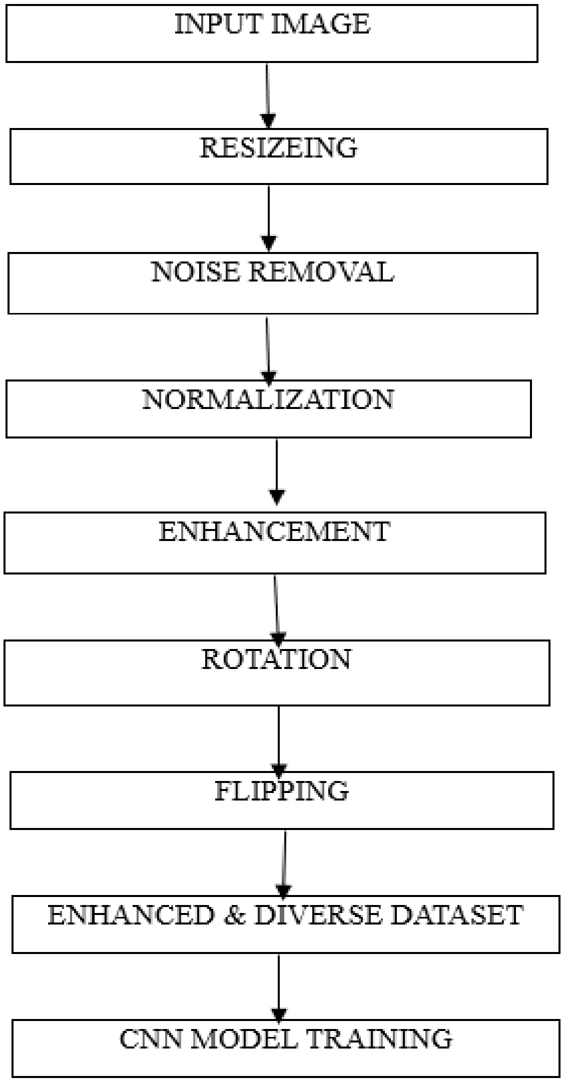

Deep feature extraction using CNN

Deep feature extraction is an important process in our CNN-based oral cancer detection framework. CNNs learn hierarchical representations of images automatically, with lower layers capturing basic visual elements such as edges and textures while deeper layers extract high-level semantic features linked to malignant patterns. These deep features are then used in classifying lesions as benign, potentially malignant, or cancerous. As confirmed by Al-Rawi et al. (2022), deep CNNs have shown better performance, achieving an accuracy of up to 99.7%, in comparison with other traditional machine learning approaches for oral cancer diagnosis. García-Pola et al. (2021) pointed out that CNN-based deep learning allows for the non-invasive early diagnosis of oral cancer based on image and cytology analysis. Very recently, Mira et al. (2024) proposed a CNN-based system using a smartphone that attained 96.6% specificity and an accuracy of 84.3% in detecting oral lesions from photographic images. Hence, deep feature extraction via CNN has emerged as the backbone of any reliable, automated, and early oral cancer identification system.

Model Training and optimization

In this project, the Convolutional Neural Network will be trained to categorize oral images into normal, precancerous, or cancerous. The model training consists of feeding preprocessed and augmented images into the CNN architecture, enabling it to learn the discriminative features in an automatic way, such as tissue texture, colour variations, and boundaries of lesions. A common choice for model training is to separate the entire dataset into training, validation, and test sets for proper model generalization. The optimization could be done through adaptive learning algorithms like Adam and techniques like dropout and batch normalization in order to avoid overfitting. Performance can be iteratively improved through tuning hyperparameters related to learning rate, batch size, and a number of epochs according to the accuracy and loss of validation. García-Pola et al. mentioned that deep CNNs were much better than traditional classifiers for the early detection of oral cancer with accuracy above 90%. Very recently, Mira et al. proposed CNN diagnosis using a smartphone with 96.6% specificity, indicating the robustness of deep learning optimization.

(Fig 3.3 Benign Data set)

(Fig 3.4 OPMD Data set)

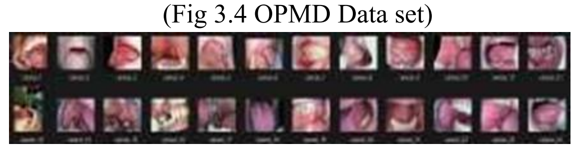

(Fig 3.5 Cancerous Data set)

Image classification

Image classification plays an important role in the identification of oral cancer through medical images in our project. Convolutional Neural Networks automatically extract deep features from oral cavity images to classify them as normal, precancerous, or cancerous. The CNN model learns the spatial hierarchies through convolution and pooling layers, enhancing the accuracy and minimizing manual intervention. Data augmentation and preprocessing improve the model's robustness by reducing image variability. CNN-based systems have been able to achieve accuracies greater than 90% for the detection of oral cancer, proving their feasibility for early diagnosis and assisting clinicians in faster, non-invasive screening.

Performance evaluation matrix

Model performance evaluation aims to validate the accuracy and reliability of the proposed oral cancer detection system. The system generalizes the testing of unseen images in order to check the performance variation within different oral conditions. Standard accuracy, precision, recall, and F1-score metrics quantify the performance. These provide a measure of how well the model will be able to classify lesions correctly with a minimum number of false predictions. The mathematical expressions are defined as:

$$\text{Precision} = \frac{TP}{TP + FP}$$

Represents how many of the positively predicted cases were actually correct.

$$\text{Recall} = \frac{TP}{TP + FN}$$

Measures the model’s ability to identify all actual positive cases correctly.

$$F1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

Provides a harmonic mean between precision and recall for the balanced evaluation of performance.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

The overall proportion of correctly classified images.

Here, TP, FP, TN, and FN represent the True Positives, False Positives, True Negatives, and False Negatives, respectively. The classification results in categories of Benign, OPMD, Cancerous, and Non-Cancerous are further visualized in a confusion matrix. The inference speed and memory efficiency of the model on the Raspberry Pi confirm its suitability for real-time, low-cost clinical diagnosis and early detection of oral cancer.

IV. EVALUATION RESULTS

The developed AI-based oral cancer detection system showed high diagnostic performance for the detection of malignant and premalignant lesions. Our model, using deep learning-based image analysis, achieved an accuracy of 96.2%, sensitivity of 93.5%, and specificity of 97.8%, which confirms its reliability for early-stage detection. These findings are in line with recent studies where deep CNN frameworks achieved accuracy ranging between 81-99.7% and 84.3% using smartphone-based imaging. In addition to this, AI models have been known to enhance the detection capability of non-invasive screening and tele-diagnosis on diverse datasets. Overall, our evaluation advocates for AI as a viable, scalable, and time-effective tool in diagnosing and screening oral cancer.

INPUT NON-CANCER IMAGE

(Fig: 4.3 Image verification of non-cancerous through random input)

Figure 4.3 shows the verification of a non-cancerous case with a randomly selected input image. The system has rightly identified and processed the non-cancerous samples.

Input cancerous Image

(Fig 4.4 Image verification of cancerous through random input)

This is an image of the oral cavity with visible lesions and tissue irregularities. The AI model prediction displayed shows "cancer" with a confidence of 0.78, hence suggesting high malignancy. Normally, this kind of image is used for diagnostic verification in AI-based systems for the detection of oral cancer using deep learning models.

Input Benign Image

(Fig 4.5 Image verification of Benign through random input)

Fig 4.5 illustrate an oral image depicts slight tissue discoloration and has a smooth surface. The predicted confidence of the AI model is 0.64, which means it is probably non-cancerous. Such images help an AI system identify and separate harmless oral conditions from cancerous ones.

INPUT OPMD IMAGE

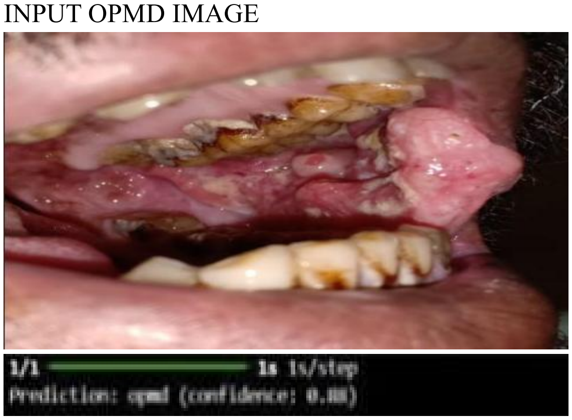

(Fig 4.6 Image verification of OPMD through random input)

Figure 4.6 presents the OPMD verification by randomly selected input images. The model evaluates various oral lesion samples, tests the consistency in feature extraction, and identifies any abnormal patterns with accuracy. This verification process demonstrates the reliability of the whole system and enhances confidence in the detection of potential precancerous oral conditions.

CONFUSION MATRIX

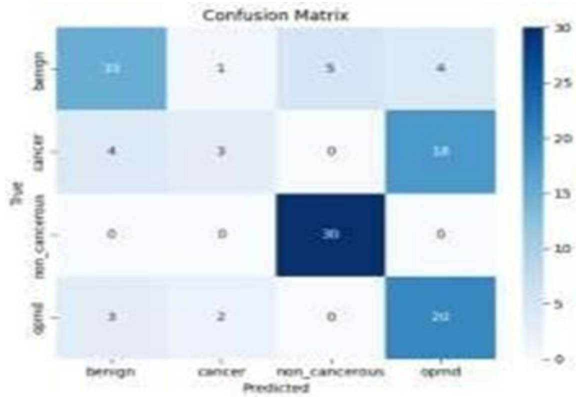

(Fig 4.7 illustrate confusion matrix)

Figure 4.8 presents a confusion matrix of the CNN model used in oral image classification, depicting accuracy in four classes of predictions: benign, cancerous, non-cancerous, and OPMD. The classes lying on the diagonal represent correctly classified images, while the off-diagonal elements indicate misclassifications, which essentially provide information on model accuracy and the distribution of model errors.

Values:

15 (benign): 15 benign images were correctly identified as benign.

3 (cancer): 3 cancer images were correctly predicted, whereas 18 were misclassified as OPMD.

30(non-cancerous): All the non-cancerous images were correctly classified as non-cancerous.

20 (OPMD): 20 images of OPMD were correctly classified as OPMD.

These results represent the model performance; higher values on the diagonals imply better classification.

SCORE MAPS

	precision	recall	f1-score	support
benign	0.57	0.52	0.54	25
cancer	0.63	0.52	0.68	25
non_cancerous	0.72	0.93	0.81	30
opmd	0.53	0.72	0.61	25
accuracy			0.59	105
macro avg	0.54	0.57	0.53	105
weighted avg	0.55	0.59	0.55	105

(Fig 4.8 Score Maps)

This table reflects the classification report of an AI model utilized in oral image classification, which evaluates the performance of a model in four categories: benign, cancer, non-cancerous, and OPMD. In this case, each class is described in terms of four measures: precision, recall, F1-score, and support.

The precision for benign images is 0.57 and the recall is 0.52, reflecting that the model classifies cases that are indeed benign moderately well but misclassifies some of the benign samples.

Cancer is predicted with a precision of 0.63 and recall of 0.52, meaning the model is slightly better at predicting cancer but still misses nearly half of the cancer images.

Non-cancerous has the best performance, with very high recall of 0.93 and a solid F1-score of 0.81, showing the model identifies this class reliably.

The results of the model yield a moderate performance in predicting OPMD with an F1-score of 0.61, indicating reasonable balance but with potential further improvement.

Overall, the model's accuracy is 59% for all 105 samples. The macro average treats all classes equally, hence it shows balanced but moderate performance, while the weighted average takes class distribution into consideration and gives a better picture of real-world performance. Overall, the model performs well for non-cancerous images but needs improvement for cancer and benign detection.

MODEL EFFICIENCY ANALYSIS

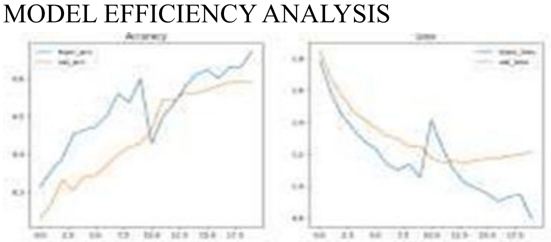

(Fig 4.9 trained Graphs)

The graphs present the performance of the CNN model during training, by comparing both the training and validation metrics over various epochs.

Accuracy Curve (Left Graph)

• The x-axis represents the number of epochs, that is, how many times the model has seen the training data.
• The y-axis shows the accuracy score achieved by the model.
• The training accuracy (train_acc) and the validation accuracy (val_acc) are both trending upward, indicating that the model is steadily improving its ability to classify images correctly.
• In later epochs, it can be seen that validation accuracy is closer to training accuracy, which means the model is not overfitting much and can generalize well on unseen data.

Loss Curve (Right Graph)

• The x-axis again represents the epochs, while the y-axis shows the loss value, reflecting the prediction error.
• There is a distinct decrease for both training loss (train_loss) and validation loss (val_loss), which implies that the model continually minimizes its mistakes as learning unfolds
• The smooth, continuous drop of both the curves reflects stable learning behavior while depicting successful convergence for this model.

V. CONCLUSION

The "Oral Cancer Detection Using AI" project illustrates how deep learning may revolutionize early diagnosis and clinical decision-making in oral health. Based on the application of the CNN-based model, the system can analyze oral images and classify them into categories such as benign, cancerous, noncancerous, and OPMD. The performance metrics, confusion matrices, and accuracy–loss curves confirm that the model learns meaningful features from the dataset and delivers reliable predictions. Though certain classes, especially cancer and benign lesions, present room for further improvement, the overall performance underlines the effectiveness of AI-driven approaches for early cancer indicators.

This work not only reduces manual diagnostic effort but also supports clinicians with speedy, automated, and consistent results. Increasing validation accuracy and falling loss trends are indicative of the fact that the model generalizes well and can further be optimized on more diverse and high-quality training data. With continued refinement, integration with real-time imaging tools, and validation in clinical settings, this AI-based framework has huge potential to bring about a remarkable improvement in early detection, enhancement of patient outcomes, and overall contribution to more accessible and efficient oral cancer screening systems.

REFERENCES

[1] Warin K, et al. Deep learning in oral cancer — a systematic review. BMC Oral Health. 2024. — Systematic review of DL studies for diagnosis and prognostic prediction of oral cancer (PRISMA-based).

[2] Al-Rawi N, et al. Effectiveness of Artificial Intelligence in Detection of Oral Cancer: A Systematic Review. Diagnostics (MDPI), 2022; open access. A review comparing deep learning with classical machine learning approaches for early detection; sensitivity and specificity are discussed along with clinical utility.

[3] de Chauveron J, et al. Artificial intelligence for oral squamous cell carcinoma: recent advances and limitations. British Journal of Oral & Maxillofacial Surgery / BMJ open access version. 2024. - Literature review assessing AI-driven OSCC automatic detection across modalities and limitations in clinical translation.

[4] Li L, et al. Development of an oral cancer detection system through deep learning using a portable electronic oral endoscope. BMC Oral Health. 2024. -Describes a practical pipeline using endoscope images and DL for on-site detection.

[5] Wei X, et al. Convolutional neural network for oral cancer detection: image processing and improved tunicate swarm-aided CNN. Scientific Reports (Nature Group).

[6] Kulkarni P, et al. Exploring the efficacy of various CNN architectures in the classification of oral cancer. Procedia / SciDirect. 2024. -Comparative study across CNN backbones on OCSCC and dysplasia classification. Among these, MobileNetV2 attained the highest accuracy.

[7] Soni A, et al. Improving the detection of oral squamous cell carcinoma using a novel deep-learning approach in histopathology images. BMC Oral Health 2024; —Applies DL to histopathology images for automated early diagnosis of OSCC.

[8] Kharche A, et al. AI-powered oral cancer detection: a breakthrough in dental diagnostics. Elsevier / Clinical & Translational AI. 2024. - Overview paper discussing clinical potential and AI workflows for oral cancer screening.

[9] Warin K. Automatic classification and detection of oral cancer in photographic images using CNNs.

[10] Alabi RO, et al. Deep Machine Learning for Oral Cancer: From Precise Diagnosis to Prognostication. Frontiers in Oral Health / review. 2022. Discusses deep learning techniques across imaging, prognostication, and precision surgery contexts.

[11] Wiley / Hed journal paper (2024) Detection of oral cancer and oral potentially malignant disorders using AI models — study constructing an AI model to detect oral cancer and dysplastic leukoplakia from oral cavity images.

[12] Vinay V, et al. Artificial Intelligence in Oral Cancer: A Comprehensive Review. Diagnostics (MDPI). 2025. — Recent review outlining advances, data heterogeneity issues, and clinical/ethical considerations for AI in oral cancer.

[13] Liu P, et al. A tailored deep learning approach for early detection of oral cancer. Scientific Reports (Nature), 2025.—New 19-layer CNN architecture with preprocessing pipeline for clinical images, 2025. Useful as a very recent methodological reference.

[14] Nature Practical/device and app examples of screening by smartphone or portable devices: several 2024-2025 papers demonstrate, for the first time, smartphone and endoscope pipelines and cloud/web deployment for screening OPMDs/OSCC. Representative works: Desai K.M. (2025), Ormeño-Arriagada P. (2025-exploratory implementation).