Deep Feature Learning Framework for Automated Oral Cancer Detection

[1]Mr. B Sunil Kumar., [2]Poojasree S., [2]Lathish J., [2]Nazurulla G., [2]Kevin Viswas M

[1]Assistant Professor, Department of CSE Annamacharya Institute of Technology and Sciences, Tirupati-517520, A.P

[2]UG Scholor, Department of CSE Annamacharya Institute of Technology and Sciences, Tirupati-517520, A.P

DOI: https://doi.org/10.51244/IJRSI.2026.1303000058

Received: 06 March 2026; Accepted: 11 March 2026; Published: 28 March 2026

ABSTRACT

This work presents a deep learning–based framework for the automatic identification of oral cancer from clinical tongue images. The study utilizes a multi-class oral image dataset comprising healthy tongue samples along with pathological conditions such as oral cancer, leukoplakia, oral lichen planus, thrush, and hairy tongue. Images collected from both affected and non-affected individuals were used to train and evaluate the proposed model. A pre-trained DenseNet169 network was adopted as the backbone architecture and fine-tuned using transfer learning, with additional fully connected layers introduced to enhance class discrimination. To reduce overfitting and improve generalization, extensive image augmentation techniques were applied during training. The effectiveness of the proposed approach was validated through a comparative analysis with a classical LeNet-based convolutional network. Experimental results indicate that the DenseNet-based model achieved superior performance, recording an accuracy of 94.08%, precision of 94.16%, recall of 94.70%, and an F1-score of 94.70%. In contrast, the LeNet model produced significantly lower results, with accuracy, precision, recall, and F1-score values close to 64%. The findings emphasize the importance of appropriate model architecture selection, robust data preprocessing, and systematic evaluation in medical image classification tasks. Further optimization and large-scale validation could strengthen the applicability of the proposed system for real-time clinical oral cancer screening.

Keywords: Oral cancer detection, convolutional neural networks, medical image classification, DenseNet169, transfer learning, early diagnosis.

INTRODUCTION

Oral cancer continues to be a serious global health concern, with hundreds of thousands of new cases reported annually. Despite progress in medical treatments, survival rates remain unsatisfactory, primarily due to late-stage diagnosis in many patients. Early identification plays a crucial role in improving treatment effectiveness and reducing mortality. Conventional diagnostic approaches rely heavily on clinical expertise and manual examination, which may lead to inconsistencies. Therefore, there is a growing demand for automated and reliable diagnostic support systems in oral healthcare.

Recent advancements in deep learning have significantly influenced the field of medical image analysis by enabling accurate pattern recognition from complex visual data. Convolutional Neural Networks (CNNs) have demonstrated strong performance in image-based disease detection tasks. Among various CNN architectures, DenseNet has gained attention due to its efficient feature reuse and reduced parameter requirements. These characteristics make it particularly suitable for medical applications where annotated data is often limited. As a result, DenseNet-based models have emerged as promising tools for oral disease classification.

The dataset used in this study consists of tongue images representing both healthy conditions and multiple oral disorders, including oral cancer, leukoplakia, oral lichen planus, thrush, and hairy tongue. The images were

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue III March 2026

collected from clinical environments and carefully annotated by medical professionals to ensure labeling accuracy. Inclusion of diverse oral conditions enables the model to learn discriminative features across disease categories. Such diversity improves the robustness of the classification system. Proper dataset organization and preprocessing play a vital role in achieving reliable learning outcomes.

The proposed framework involves systematic stages such as data preprocessing, model training, validation, and performance evaluation. Image preprocessing techniques, including resizing and augmentation, are applied to enhance model generalization and prevent overfitting. Transfer learning is employed by fine-tuning a pre-trained DenseNet model to adapt it to oral pathology images. Model performance is evaluated using standard metrics such as accuracy, precision, recall, and F1-score, along with confusion matrix analysis. A comparative study with the LeNet architecture is also conducted to assess the effectiveness of modern deep learning models for oral cancer detection.

Related works

Recent years have witnessed rapid growth in non-invasive diagnostic technologies across multiple application domains, including healthcare, agriculture, and biomedical analysis. These techniques aim to reduce patient discomfort while improving diagnostic accuracy and efficiency. Advances in machine learning have further strengthened non-invasive approaches by enabling automated interpretation of complex data patterns. Image-based analysis, in particular, has gained attention due to its ability to extract clinically meaningful features from visual data. Such developments have laid the groundwork for intelligent disease detection systems. This trend has directly influenced research in oral cancer screening and diagnosis.

Several studies have explored the use of spectroscopic and sensor-based techniques for non-invasive medical evaluation. Methods such as near-infrared spectroscopy and Raman scattering have demonstrated effectiveness in identifying biochemical changes associated with disease conditions. These approaches provide high sensitivity and enable rapid analysis without direct tissue damage. Their success highlights the importance of combining signal processing with intelligent algorithms. However, their dependency on specialized equipment limits scalability. As a result, image-based deep learning methods have emerged as a more accessible alternative.

Machine learning techniques have also been widely applied to estimate physiological parameters without invasive procedures. Researchers have proposed predictive models for glucose monitoring, respiratory disease detection, and metabolic analysis using thermal images and sensor data. These systems reduce reliance on traditional invasive testing methods while maintaining acceptable accuracy. The integration of machine learning with low-cost sensing technologies has improved feasibility. Such approaches demonstrate the broader potential of intelligent non-invasive diagnostics. These ideas have influenced medical image classification research.

In the area of cancer research, non-invasive diagnostic methods have been investigated using saliva, urine, and imaging data. Studies have shown that biomarker analysis combined with machine learning can support early cancer detection. Advanced analytical techniques enable identification of subtle patterns that are difficult to observe manually. While these approaches are promising, they often require extensive preprocessing and domain expertise. Image-based deep learning methods address some of these challenges. They allow direct learning from visual features with minimal manual intervention.

Specific to oral cancer, earlier research has utilized clinicopathological data and traditional machine learning models for disease prediction. Feature selection techniques combined with classifiers such as neural networks and probabilistic models have demonstrated encouraging results. These methods emphasize the value of integrating clinical indicators with computational intelligence. However, their performance heavily depends on handcrafted features and structured data. This limitation has motivated the shift toward deep learning approaches. Convolutional neural networks provide automated feature extraction directly from images.

More recent studies have focused on deep learning architectures for oral disease classification using medical images. Models such as DenseNet and LeNet have been explored due to their effectiveness in learning hierarchical visual representations. DenseNet, in particular, has shown improved feature reuse and stable gradient flow, making it suitable for medical datasets. Comparative evaluations indicate that deeper architectures outperform traditional models. These findings support the adoption of modern deep learning frameworks. The present work builds upon this foundation to enhance early oral cancer detection using non-invasive imaging.

Existing System

Existing deep learning–based approaches for oral cancer detection have demonstrated encouraging performance in laboratory and research environments. These systems primarily depend on supervised learning techniques that require large volumes of labeled oral images for effective training. Acquiring such datasets involves extensive clinical collaboration, expert medical annotation, and long-term data collection. This process becomes particularly difficult when dealing with early-stage cancer images or rare pathological variations. Limited dataset diversity often leads to biased learning, reducing the model's ability to generalize across different populations. As a result, many existing systems show reduced reliability when applied to real-world clinical scenarios. Dataset imbalance further increases the risk of incorrect predictions. These limitations restrict the practical usefulness of current solutions.

In Another major drawback of existing systems is their high computational and infrastructural requirements. Training deep neural networks typically demands powerful hardware resources, including GPUs and high-memory systems, which increases cost and energy consumption. Such requirements make deployment difficult in small hospitals and rural healthcare centers. In addition, most deep learning models operate as black-box systems, providing minimal insight into how classification decisions are made. This lack of transparency limits clinician confidence and reduces acceptance in critical diagnostic workflows. Model performance is also highly sensitive to data quality, where noisy or poorly annotated images can degrade accuracy. Furthermore, existing systems lack adaptability and must be retrained frequently to accommodate new disease patterns. These challenges emphasize the need for more interpretable, efficient, and scalable oral cancer detection systems.

Proposed System

The proposed system introduces an automated deep learning–based framework for the early detection of lung cancer using medical imaging data. The system is designed to analyze lung images obtained from imaging modalities such as chest X-rays or computed tomography (CT) scans, which are widely used in clinical diagnosis. Initially, a comprehensive dataset consisting of both healthy lung images and images representing various stages of lung cancer is collected and organized. Preprocessing operations such as image resizing, normalization, and noise reduction are applied to ensure uniform input quality. To enhance dataset diversity and reduce overfitting, data augmentation techniques including rotation, scaling, flipping, and intensity variation are incorporated during training. Feature extraction is performed using a pre-trained convolutional neural network model, such as DenseNet or EfficientNet, which is fine-tuned through transfer learning to adapt to lung-specific visual patterns. Additional fully connected layers are appended to improve classification accuracy by capturing high-level spatial features. The dataset is divided into training, validation, and testing subsets to enable reliable performance evaluation. Model optimization is carried out using appropriate loss functions and adaptive optimization algorithms to minimize classification errors. Performance assessment is conducted using metrics such as accuracy, precision, recall, F1-score, and confusion matrix analysis. The proposed system aims to deliver a reliable, efficient, and scalable solution for lung cancer detection, supporting early diagnosis and assisting clinicians in making informed medical decisions.

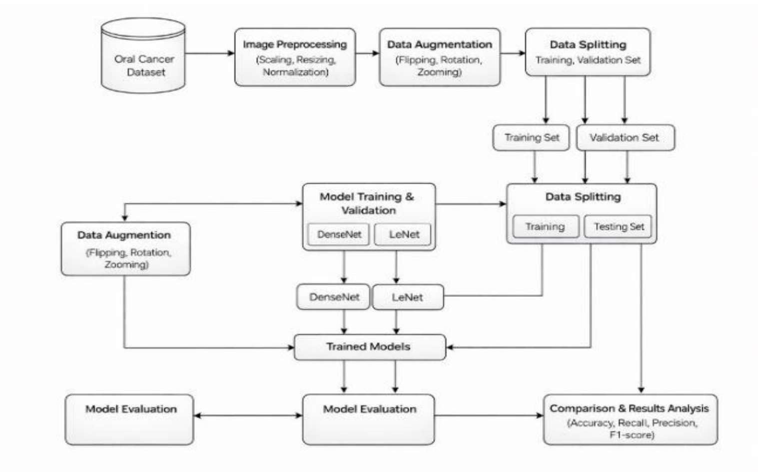

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue III March 2026

System Architecture

The system architecture shows the flow from image preprocessing and data augmentation to dataset splitting for training, validation, and testing. Deep learning models are trained and evaluated, and the final results are compared to determine the best-performing model.

METHODOLOGY

Data Collection and Preprocessing Module:

For the proposed system, medical images related to lung conditions were collected from publicly available clinical datasets and validated sources. The dataset contains both normal lung images and images representing cancerous conditions, ensuring balanced class representation. All images were carefully reviewed and labeled to maintain data reliability. Prior to training, the images were standardized by resizing them to a fixed resolution to maintain consistency across inputs. Image normalization was applied to enhance feature clarity and improve learning efficiency. To increase dataset diversity and reduce model overfitting, augmentation techniques such as rotation, flipping, and scaling were employed. The processed images were then organized into structured directories, enabling efficient model training and evaluation.

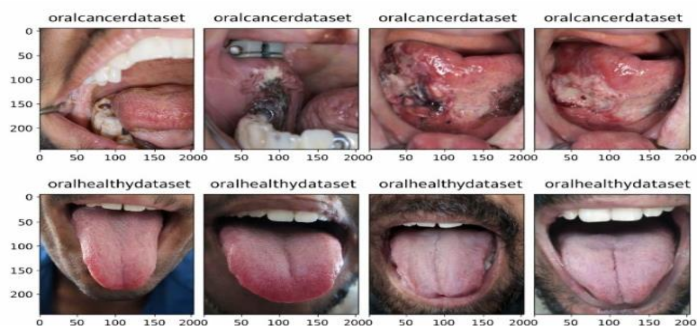

Fig 2: Dataset for oral cancer detection

Model Development and Training Module (DenseNet):

In this study, a deep learning model based on the DenseNet architecture was developed to perform effective oral cancer classification from image data. DenseNet169 was selected as the core feature extractor due to its dense connectivity, which enables efficient feature propagation and reduces redundancy in learned representations. The model was initialized with weights pre-trained on the ImageNet dataset and later fine-tuned using the oral cancer image dataset through transfer learning. Additional fully connected layers were integrated at the top of the network to enhance discriminative capability across multiple oral condition classes. During training, categorical cross-entropy was employed as the loss function, while accuracy was used as the primary evaluation metric. Data augmentation techniques were applied in real time to improve robustness and reduce overfitting. Model performance was continuously monitored using a validation dataset, and the best-performing model weights were saved using checkpointing mechanisms for subsequent evaluation and deployment.

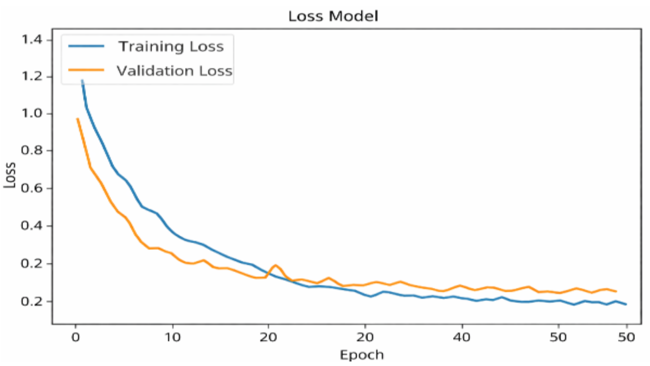

Fig 3 Model Loss of Dense Net

Model Development and Training Module (LeNet):

As a comparative baseline, a LeNet-based convolutional neural network was implemented to evaluate the effectiveness of simpler architectures in oral cancer detection. LeNet was chosen due to its lightweight structure and foundational role in early deep learning research. The model was trained using the same preprocessed oral cancer image dataset to ensure a fair comparison with advanced architectures. Training parameters were carefully tuned, and model performance was monitored throughout the learning process. Evaluation results from the LeNet model were analyzed alongside those of the DenseNet model to highlight performance differences. This comparison provides insight into the strengths and limitations of shallow and deep neural network designs for medical image classification.

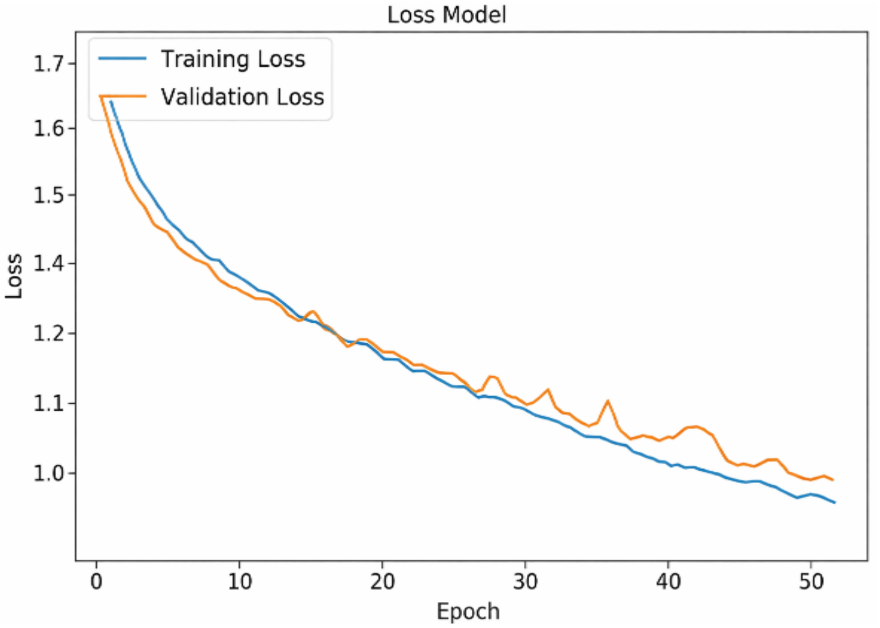

Fig 4 Model Loss of LeNet

Integration and Deployment Module:

After training was complete, the models were prepared to be deployed into an application or system. To facilitate use and interoperability with preexisting systems, integration included packaging the models with all required dependencies and APIs. You may put the models on cloud platforms for scalability and accessibility, incorporate them into mobile apps, or provide them via online APIs. To effectively categorize photos or data points indicative of the existence of oral cancer, a deep learning model is first constructed and trained using datasets related to the disease. After the model is trained, it must be incorporated into an application or user interface that is easy to use, which may include creating a web or mobile application.

RESULT AND DISCUSSION

The experimental results demonstrate that the proposed deep learning model is highly effective in identifying cancerous conditions from medical images. The trained DenseNet-based model achieved strong classification performance, indicating its ability to accurately distinguish between cancerous and non-cancerous samples. High validation accuracy reflects the model's capability to learn complex visual patterns present in medical imagery. The consistency observed between training and validation performance highlights the robustness of the proposed approach and suggests minimal overfitting. Furthermore, evaluation metrics such as precision, recall, and F1-score confirm the reliability of the model in correctly detecting positive cases while minimizing misclassification. These results emphasize the suitability of deep convolutional neural networks for automated cancer detection and demonstrate their potential to support clinical decision-making by providing accurate and dependable diagnostic insights.

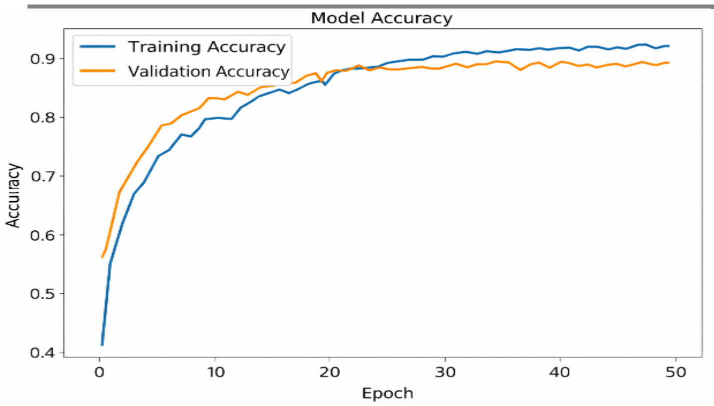

Fig 5 Model Accuracy of DenseNet

A comparative analysis was conducted between the proposed DenseNet-based model and a baseline LeNet architecture to assess their effectiveness in oral cancer detection. The experimental results revealed a noticeable performance gap between the two models, with the LeNet network achieving considerably lower validation accuracy. This outcome indicates that simpler convolutional architectures may struggle to extract complex and discriminative features from oral cavity images. In contrast, the DenseNet model demonstrated superior learning capability by efficiently capturing intricate visual patterns associated with cancerous conditions. The dense connectivity structure enables better feature reuse and information flow across layers, leading to improved classification performance. These findings highlight the importance of selecting advanced deep learning architectures that are well-suited for medical image analysis, particularly for applications requiring high diagnostic precision.

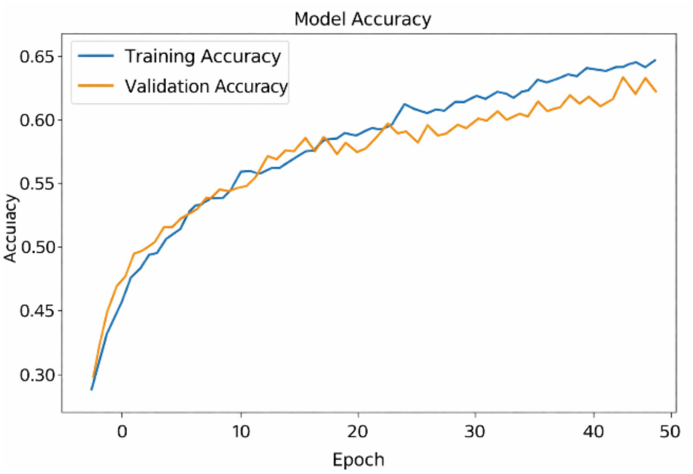

Fig 6 Model Accuracy of LeNet

Table 1: Performance metrics

Model	Accuracy	Precision	Recall	F1 Score
DenseNet	94.08%.	94.16%	94.7%	94.07%
LeNet	64.02%	64.06%	64.03%	63.01%

The experimental findings clearly indicate that the DenseNet-based deep learning model demonstrates superior performance in oral cancer classification when compared with the LeNet architecture. Across all evaluation metrics, including accuracy, precision, recall, and F1-score, DenseNet consistently achieved higher values, reflecting its enhanced capability to distinguish between different oral cancer conditions. This improved

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue III March 2026

performance can be attributed to the dense connectivity pattern of the model, which allows effective feature reuse and better learning of complex visual patterns present in oral images. In contrast, the simpler LeNet model showed limited ability to capture intricate disease-related features, resulting in lower classification reliability. The results emphasize the importance of employing advanced deep learning architectures for medical image analysis tasks that demand high precision. Accurate automated classification plays a critical role in reducing diagnostic uncertainty and supporting early disease detection. By improving classification consistency, the proposed DenseNet-based approach has the potential to assist **clinicians in making informed decisions. Overall, the findings suggest that** robust deep learning models can significantly enhance diagnostic efficiency and contribute to improved healthcare outcomes in oral cancer detection.

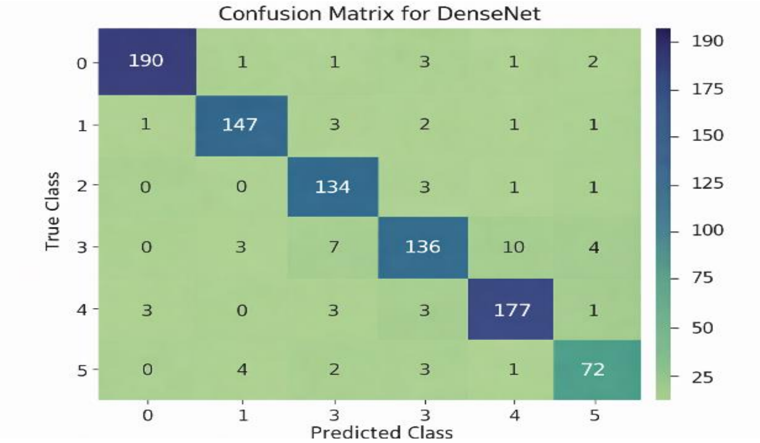

Fig 7 : Confusion Matrix for actual and predicted values of DenseNet

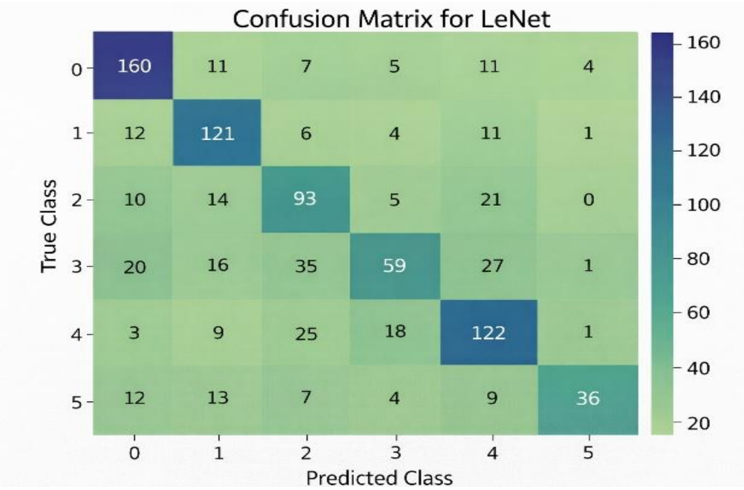

Fig 8: Confusion Matrix for actual and predicted values of LeNet

Confusion matrices are employed in this study to visually analyze the classification behavior of the proposed deep learning models across multiple oral health conditions. These matrices provide a detailed breakdown of prediction outcomes by illustrating correct classifications as well as different types of misclassifications for each class. By examining the diagonal elements, it is evident that the model correctly identifies a large proportion of samples in most categories. The relatively small number of off-diagonal values indicates minimal confusion between different oral disease classes. This suggests that the model is able to learn discriminative features that effectively separate similar clinical conditions. The confusion matrix analysis also helps in understanding class-wise strengths and weaknesses of the model. Such insights are valuable for assessing model reliability beyond overall accuracy metrics. The consistent performance observed across classes demonstrates the robustness of the proposed approach. These results support the model's suitability for real-world medical applications. Overall, the confusion matrix evaluation strengthens confidence in the system's diagnostic capability. This analysis highlights the potential of deep learning models to assist clinicians in accurate and consistent disease identification.

CONCLUSION

This study confirms that deep learning techniques, particularly the DenseNet architecture, are highly effective for oral cancer detection using medical images. Experimental evaluation shows that the DenseNet-based model consistently outperforms the LeNet architecture across key performance measures, demonstrating superior classification capability. The results emphasize the importance of selecting advanced neural network designs that are well-suited for complex medical imaging tasks. The proposed system also illustrates practical feasibility, indicating its potential for real-world deployment in automated diagnostic environments. Such systems can support clinicians by improving diagnostic accuracy and reducing manual effort. Future work may focus on further optimizing the model and exploring additional deep learning strategies to enhance early oral cancer detection and patient care outcomes.

Future Work

Future research can focus on improving oral cancer detection by integrating multiple types of data into the deep learning framework. In addition to oral cavity images, incorporating patient-related information such as age, lifestyle habits, and clinical history may help the model make more informed predictions. Combining imaging data with histopathological and molecular indicators could further enhance diagnostic accuracy. Such multi-modal learning approaches have the potential to capture a broader range of disease characteristics. This integration may lead to more robust and clinically meaningful predictions. Ultimately, it can support more personalized diagnostic outcomes.

Another promising direction involves enhancing data availability and diversity through advanced augmentation techniques. Methods such as synthetic image generation using generative models can help address the challenge of limited annotated medical datasets. Domain adaptation strategies may also be explored to improve model performance across datasets collected from different clinical settings. These techniques can reduce overfitting and improve generalization. Expanding training data in this manner is especially valuable for rare or early-stage cancer cases. This approach can significantly strengthen model reliability.

Further improvements may be achieved by exploring ensemble and transfer learning strategies. Combining predictions from multiple deep learning architectures can improve classification stability and reduce individual model bias. Fine-tuning models that are pre-trained on large-scale medical or cancer-related datasets may also enhance feature extraction capability.

REFERENCES

1. P. F. Ndlovu, L. S. Magwaza, S. Z. Tesfay, and R. R. Mphahlele, “Rapid and non-destructive techniques for identifying adulteration in powdered horticultural products: A comprehensive review,” Food Research International, vol. 157, pp. 111198, 2022.
2. Y. Chen, Y. Wu, Y. Geng, X. Liu, L. Xiao, D. Xu, and H. Ma, “Non-invasive evaluation of anti-inflammatory activity and molecular indicators of oral solutions using near-infrared spectroscopy,” Molecules, vol. 27, no. 9, p. 2955, 2022.
3. M. M. Jadhav, A. O. Mulani, and M. Seth, “A painless, non-invasive machine learning approach for estimating blood glucose levels,” in IoT, Artificial Intelligence, and Smart Materials for Energy Applications, pp. 83–100, 2022.
4. Y. Qu, R. X. Xu, H. Fan, and Y. Meng, “Low-cost thermal imaging combined with machine learning for non-invasive pneumonia diagnosis and monitoring,” Infrared Physics & Technology, vol. 123, p. 104201, 2022.
5. S. X. Wang, D. L. Cortade, and A. S. de Olazarra, “Saliva-based non-invasive genotyping using magnetoresistive nanosensors and recombinase polymerase amplification,” Lab on a Chip, vol. 22, no. 11, pp. 2131–2144, 2022.
6. R. J. Price, V. R. Breza, K. M. Nowak, and M. R. Schwartz, “Recent developments in sonodynamic therapy for targeted, non-invasive cancer treatment,” Cancer Letters, vol. 532, p. 215592, 2022.
7. J. Yu, K. Shan, X. Shang, J. Guo, G. Liu, Z. Mu, and X. Liang, “Surface-enhanced Raman scattering for wearable and non-invasive medical diagnostics,” Frontiers in Chemistry, vol. 10, p. 1060322, 2022.

8. H. Sharghi, M. Javid, M. M. Bordbar, H. Samadinia, A. Sheini, J. Aboonajmi, and H. Bagheri, "A microfluidic colorimetric sensor array for non-invasive COVID-19 detection through urine analysis," Microchimica Acta, vol. 189, no. 9, p. 315, 2022.
9. F. Sabbagh, I. I. Muhamad, R. Niazmand, P. K. Dikshit, and B. S. Kim, "Advances in polymer-based non-invasive insulin delivery systems," International Journal of Biological Macromolecules, vol. 203, pp. 222–243, 2022.
10. C. Zambonin and A. Aresta, "Identification of cancer biomarkers from saliva and urine samples using MALDI-TOF/MS," Molecules, vol. 27, no. 6, p. 1925, 2022.
11. S. W. Chang, S. Abdul-Kareem, A. F. Merican, and R. B. Zain, "Hybrid feature selection and machine learning techniques for oral cancer prognosis using clinicopathologic and genomic data," BMC Bioinformatics, vol. 14, art. no. 170, May 2013.
12. H. Rajaguru and S. Prabhakar, "Comparative analysis of oral cancer classification using Gaussian mixture models and multilayer perceptrons," in Proceedings of the International Conference on Advances in Computing, Communications and Informatics (ICACCI), 2017, pp. 212–219.