Automated Oral Cancer Identification: Deep Learning Integration with AI-Enhanced Diagnostic Pathways

[1]Amulya S S, [2]Dr. Nanditha B R, [3]H M Chinthana, [4]Amulya Patel, [5]Gowri S P

[1]*Information and Science and Engineering, Malnad College of Engineering* Hassan-573202, India

[2]*Associate Professor, Information and Science and Engineering, Malnad College of Engineering* Hassan-573202, India

[3]*Information and Science and Engineering, Malnad College of Engineering* Hassan-573202, India

[4]*Information and Science and Engineering, Malnad College of Engineering* Hassan-573202, India

[5]*Information and Science and Engineering, Malnad College of Engineering* Hassan-573202, India

Email id: amulyasrinivas86@gmail.com, brn@mcehassan.ac.in, hmchinthana21@gmail.com ,amulyapatel183@gmail.com, gowri03sp@gmail.com

***Abstract*— Oral cavity cancer is a growing health concern, particularly in areas where the consumption of tobacco and other harmful substances is high. This disease is often diagnosed at a later stage, making treatment more difficult and increasing the mortality rate. Artificial intelligence, especially deep learning-based image classification, offers an effective approach to identifying cancerous regions at an early stage. Therefore, our study proposes to analyze and implement AI techniques to classify intraoral images into two categories: Cancerous or Non-cancerous. Early detection allows for timely medical intervention, reducing complications and improving survival outcomes. In this research, we developed and trained deep learning models using Convolutional Neural Networks (CNNs), focusing on binary classification. We compared the performance of two architectures — VGG19 and Dense Net — on a curated dataset of oral images. Among these, the VGG19 model showed superior accuracy in prediction. The project demonstrates how AI and deep learning can be leveraged to build low-cost, scalable, and accurate diagnostic tools that can potentially save lives by enabling timely intervention. With further enhancement and clinical validation, this model can be deployed in hospitals, dental clinics, and mobile health units for mass screening of oral cavity cancer.**

***Keywords*—** *Machine learning, Deep learning, Oral Cavity Cancer diagnosis, Image processing, Convolutional neural networks, Medical imaging.*

I. INTRODUCTION

Particularly in areas like Karnataka and North India, where socioeconomic factors, high tobacco use, and restricted access to healthcare worsen the disease's prevalence, oral cancer is a serious public health concern. According to studies, these areas have a huge influence on India's oral cancer burden, which makes up approximately one-third of all cases worldwide. Poor prognoses are frequently the result of late-stage detection, which there is still a major challenge despite advancements in diagnostic technologies.

To address them diagnostic challenges, this paper investigates the application of deep learning (DL) and artificial intelligence (AI) techniques. The main goal is to synthesize findings from a review of 5 research papers, covering approaches examples include support vector machines (SVMs), convolutional neural networks (CNNs), and sophisticated image processing techniques. These studies emphasize scalability and real-world applicability while highlighting the potential of integrating AI for early diagnosis classification of Oral Lesions into cancerous and non-cancerous categories.

A thorough comparison of these approaches is given in the sections that follow, highlighting important developments, gaps, and trends. The goal of the paper is to open the door for more reliable, scalable, and easily accessible AI-driven diagnostic tools by combining results from various methodologies, especially environments with restricted resources, such as rural India.

This study proposes a comparative analysis of VGG19 and DenseNet121 for the classification of oral cancer images. The primary goal is to assess the effectiveness of these models in distinguishing between normal and cancerous oral tissues, with a focus on accuracy, precision, recall, and F1-score. Our experiments reveal that DenseNet121 consistently outperforms VGG19, achieving an accuracy in the range of 96–97%, thereby demonstrating its superior feature extraction capability.

The structure of this research paper is organized to provide a comprehensive understanding of the proposed oral cancer detection approach. The paper begins with an Introduction that outlines the motivation behind the study, the challenges associated with traditional diagnostic techniques, and the potential of deep learning in medical image classification. This is followed by the Related Work section, which reviews existing studies on oral cancer detection and highlights the performance and limitations of previously used CNN models. The Methodology section details the dataset used, preprocessing steps, model architecture configurations for VGG19 and DenseNet121, and the experimental setup employed for training and testing. In the Results and Discussion section, the outcomes of the experiments are presented, including a comparative analysis based on accuracy, precision, recall, and F1-score, with a focus on the superior performance of DenseNet121. Finally, the paper concludes with the Conclusion and Future Work section, summarizing key findings and proposing possible directions for future enhancements, such as real-time deployment and testing with larger, more diverse datasets.

II. RELATED WORKS

A. *Paper title: Convolutional Neural Network-Based Clinical Predictors of Oral Dysplasia: Class Activation Map Analysis of Deep Learning Results, Authors: Seda Camalan, Hanya Mahmood, Hamidullah Binol, Anna Luiza Damaceno Araújo, Alan Roger Santos-Silva, Pablo Agustin Vargas, Marcio Ajudarte Lopes, Syed Ali Khurram, Metin N. Gurcan, Year of publication : 2021*

Description: Convolutional neural networks (CNNs) and class activation maps (CAMs) are used in this paper's deep learning-based system to categorize oral lesion images as "normal" or "suspicious," with the goal of enhancing the early identification of oral dysplasia and cancer. With pre-trained models such as Inception-ResNet-v2, the study uses transfer learning to achieve up to 95.2% accuracy on cross-validation and 86.5% accuracy on independent datasets. The system highlights crucial areas affecting predictions with interpretable heatmaps and employs image patches for improved performance. The method shows promise for automated, interpretable diagnostic tools in oral healthcare, even with a small dataset.

Methodology: Inception-ResNet-v2, ResNet-101, VGG-16, and Inception-v3 are pre-trained models that are used in this paper's CNN-based transfer learning method to categorize oral lesion images as "normal" or "suspicious." To deal with the small dataset size, image preprocessing entails splitting annotated images into smaller patches and then data augmentation. Class activation maps (CAMs) are created to visualize the regions influencing the classification, and the models are refined to predict lesion classes. The accuracy and generalizability belonging to the system are evaluated through independent testing and cross-validation on datasets from the UK and Brazil.

Limitations: A small dataset, manual lesion segmentation, dependence on photographic images without additional modalities, and dataset variability are some of the limitations of the paper. It also has interpretability issues with class activation maps (CAMs) and does not integrate clinical metadata. The goal of future research is to close these gaps.

Key Insights: The study shows how CNNs and class activation maps (CAMs) can efficiently categorize oral lesions in order to detect dysplasia early. Performance is enhanced by patch-based analysis, and generalizability is guaranteed by validation on datasets from various geographical locations. CAMs improve the interpretability of the model, making it a potentially useful AI-powered diagnostic tool.

B. *Paper title: Automated Detection and Classification of Oral Lesions Using Deep Learning for Early Detection of Oral Cancer, Authors: Roshan Alex Welikala, Paolo Remagnino, Jian Han Lim, Chee Seng Chan, Senthilmani Rajendran, Thomas George Kallarakkal, Rosnah Binti Zain, Ruwan Duminda Jayasinghe, Jyotsna Rimal, Alexander Ross Kerr, Rahmi Amtha, Karthikeya Patil, Wanninayake Mudiyanselage Tilakaratne, John Gibson, Sok Ching Cheong, Sarah Ann Barman, Year of publication : 2020*

Description: The goal of the paper, " Automated Detection and Classification of Oral Lesions Using Deep Learning for Early Diagnosis of Oral Cancer," is to improve oral cancer early detection, especially in low- and middle-income nations, by utilizing deep learning techniques. The study uses Faster R-CNN for detection of objects and ResNet-101 Image categorization is a component of the *MeMoSA® (Mobile Mouth Screening Anywhere) project. The work demonstrates the promise of AI to address global health issues like oral cancer through introducing novel strategies for combining clinician annotations and achieving encouraging initial results in lesion identification and referral decisions.

Methodology: In order to facilitate an early diagnosis of oral cancer, especially in areas with limited resources, this work proposes a deep learning-based method for the automatic identification and categorization of oral lesions. A dataset of 2,155 photos of the oral cavity was gathered and interpreted as part of the MeMoSA® (Mobile Mouth Screening Anywhere) project utilizing a innovative approach that incorporates feedback from several doctors. Two methods were evaluated: object detection using Faster R-CNN, which showed the F1 score of 41.18% for finding lesions that required referral, and picture classification using ResNet-101, which achieved an F1 score of 87.07% for recognizing the existence of lesions. The approach shows promise for scalable, affordable oral cancer screening systems by combining data enhancement, transfer learning, and performance evaluation across binary and multi-class tasks.

Limitations: A key paper's shortcomings is its very small and diverse dataset, which has an impact on the generalizability and effectiveness of the model. Furthermore, A single physician annotated a section of the dataset, which compromised the credibility of the composite annotation. Low-resolution photos are also featured in the dataset, which affects the accuracy of the model. Last but not least, the object detection method performed worse than picture classification due to its difficulties with precise lesion localization.

Key Insights: With ResNet-101 obtaining excellent accuracy in lesion identification and referral categorization and Faster R-CNN providing focused lesion localization despite its difficulties, this work shows the potential of deep learning for early oral cancer detection. Clinician inputs are combined in a novel composite annotation technique to provide trustworthy training data. The findings imply that to increase performance and generalizability, more, better-quality datasets are required. This work shows potential as a low-cost, scalable screening option for resource-constrained environments is part of the MeMoSA® platform.

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 05 | May - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

C. *Paper title: Fine-Tuning Deep Learning Architectures for Early Detection of Oral Cancer Authors: Roshan Alex Welikala, Paolo Remagnino, Jian Han Lim, Chee Seng Chan, Senthilmani Rajendran, Thomas George Kallarakkal, Rosnah Binti Zain, Ruwan Duminda Jayasinghe, Jyotsna Rimal, Alexander Ross Kerr, Rahmi Amtha, Karthikeya Patil, Wanninayake Mudiyanselage Tilakaratne, John Gibson, Sok Ching Cheong, Sarah Ann Barman , Year of publication : 2020*

Description: The focus is on binary image categorization "referral" vs. "non-referral" cases, this paper investigates the adoption with the deep learning algorithms for early-stage oral cancer detection is part of the MeMoSA® (Mobile Mouth Screening Anywhere) project. Using transfer learning on a small dataset of photographs of the oral cavity the authors compare a numberous pre-trained convolutional neural network (CNN) architectures (VGG-16, VGG-19, Inception-V3, ResNet-50, ResNet-101). Using VGG-19, the study obtains an accuracy of 80.88%, sensitivity of 85.71% and specificity of 76.42%. The findings show how AI in telemedicine has the potential to improve the accuracy of early cancer detection and referrals, especially in low- and middle-income nations.

Methodology: The MeMoSA® project dataset is utilized in the research to categorize oral cavity photographs as "referral" or "non-referral" through early detection using deep learning oral cancer detection. Transfer learning was used to refine five pre-trained CNN architectures (VGG-16, VGG-19, Inception-V3, ResNet-50, and ResNet-101). The initial layers were frozen to preserve general features, and the subsequent layers were modified for the task. Small dataset issues were resolved by batch normalization adjustments, and images were resized and enhanced to increase model robustness. Stochastic gradient descent with a weighted loss function was used in training to rectify the imbalance between classes. With sensitivity (85.71%) and specificity (76.42%), VGG-19 had the best accuracy (80.88%). Particularly in low-resource environments, the study shows how AI in telemedicine can improve early oral cancer diagnosis.

Limitations: The main drawback of the study is the small dataset size (2,155 images), which makes it difficult to train custom architectures or completely fine-tune the deep learning models, which could result in overfitting. Deeper networks like Inception-V3 and ResNet performed worse as a result of batch normalization problems brought on by domain differences between the pre-trained ImageNet data and the oral cancer dataset. The capacity to capture domain-specific features may also be restricted by the use of pre-trained models from unrelated domains. To promote clinical confidence in AI-driven decisions, future research calls for a bigger dataset, domain-specific pre-training, and more investigation into interpretability.

Key Insights: This study focuses on binary image classification of "referral" vs. "non-referral" cases to demonstrate how deep learning can enhance early oral cancer detection through telemedicine. VGG-19 outperformed the five pre-trained CNN architectures that were tested, achieving 80.88% accuracy, 85.71% sensitivity, and 76.42% specificity. In order to adapt these models to a small, domain-specific dataset, transfer learning also fine-tuning were essential. Overfitting and domain shift issues were addressed by data augmentation and batch normalization adjustments. The study emphasizes how AI-driven solutions, like the MeMoSA® project, offer the possibility for improvement diagnostic accessibility and accuracy, especially in low- and middle-income countries where a lack of specialists leads to delayed diagnoses and worse outcomes.

D. *Paper title: Deep Learning-based Cavity Detection in Diverse Intraoral Images: A Web-based Tool for Accessible Dental Care, Authors: Siji Rani S, Srija Garine, Papolu Hema Janardhana, Lakkireddy Lakshmi Prabhanjan Reddy, Penubothu Jagadeesh Venkata Kumar, Chapa Gagan Dwaz, Year of publication : 2024*

Description: This study introduces "CatchCavity," a web-based tool that uses deep learning models to identify dental cavities in a variety of intraoral photos taken with mobile cameras. The system provides an easy-to-use diagnostic tool for evaluating oral health by analyzing images for cavities using convolutional neural networks (CNNs). The study highlights the possibilities for cutting-edge AI technologies to enhance oral healthcare accessibility and diagnostic precision, particularly for users without access to specialized dental equipment, with features like secure patient record storage, performance metrics like 98.7% accuracy, and integration of location-based services for nearby dental facilities.

Methodology: Employing neural networks with convolutions (CNNs) like ResNet-50, VGG16, MobileNet, ShuffleNet, and NASNetMobile, the study uses a deep learning-based cavity detection methodology. To overcome dataset limitations, the method starts with data preprocessing, which includes augmentation techniques. In order to avoid overfitting, the models are adjusted and modified for binary classification using regularization. The system incorporates image preprocessing techniques such as noise reduction, contrast enhancement, and resizing. A Flask-built web application manages report generation, dental image uploads, and user authentication. The approach achieves high diagnostic accuracy for mobile-captured dental images by emphasizing reliable model development, validation, and deployment.

Limitations: The study's main drawback is the small dataset size, which makes it difficult to achieve reliable generalization and raises the possibility of overfitting. Furthermore, using intraoral photos taken with a mobile device introduces noise and distortion that can change image quality, which could have an effect on model performance. Although the system appears promising, access to a larger, more varied dataset could greatly increase its accuracy. Furthermore, complicated dental conditions that call for more nuanced diagnostics may be missed if binary classification is used.

Key Insights: The study demonstrates the revolutionary potential of deep learning within dental diagnostics by showing that models like ShuffleNet and VGG16 can detect cavities from intraoral images taken with a mobile device with a top accuracy of 98.7%. It emphasizes how crucial data augmentation and preprocessing are to improving model performance and lowering overfitting. The study also shows

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 05 | May - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

how dental care can be made more accessible, particularly for those without access to specialized equipment, by incorporating a web-based tool with features like location-based recommendations and secure record storage.

E. *Paper title: Detection of Cavities from Oral Images using Convolutional Neural Networks, Authors: Mohammed Abdul Hafeez Khan, Giri Prasad S., J. Angel Arul Jothi, Year of publication : 2022*

Description: A new CNN model named Dental-Net is provided in the document "Detection of Cavities from Oral Images using Convolutional Neural Networks" for automated cavity detection from oral photographic images. Through data augmentation and grayscale normalization, the study offers a solid solution to medical imaging problems brought passed by a variety of image sources. Dental-Net achieved a great precision of 94.25% on training sets and 91.09% on validation sets after being trained on a dataset of 609 images. In terms of accuracy and robustness, the model outperformed cutting-edge pre-trained models such as VGG16, MobileNetV2, InceptionV3, and ResNet18. With the goal of developing effective, scalable deep learning solutions for dental health diagnostics, this work shows great promise for lowering diagnostic complexity and costs.

Methodology: The process entails creating a brand-new CNN model called Dental-Net to identify dental cavities in photos of the mouth. 609 images were gathered from numerous sources, such as repositories and web scraping. To handle illumination variability, preprocessing techniques such as resizing, grayscale conversion, and normalization were applied to the images. To improve the robustness of the model, data augmentation methods like random rotation, zoom, and brightness adjustments were used. Four convolutional layers with max-pooling, batch normalization, and dropout layers make up Dental-Net's architecture, which ends with a dense layer for binary classification. Using binary cross-entropy loss and the Adam optimizer, the model was trained for 25 epochs, attaining high accuracy without overfitting. When performance was in contrast to pre-trained models like VGG16 and ResNet18, it showed better generalization and accuracy.

Limitations: The study points out a few drawbacks with the suggested Dental-Net model. With only 609 photos, the dataset is comparatively small, which might restrict the model's applicability to larger populations and more varied circumstances. Because the photos came from different platforms, there may have been biases introduced by variations in lighting and image quality. Furthermore, the model's dependence on grayscale images might cause it to lose crucial color information that might be useful for cavity detection. Although Dental-Net performed better than pre-trained models, more improvement is required before it can be used for real world. To guarantee dependability and scalability in real-world applications, the model still needs more testing with bigger datasets, better image acquisition methods, and integration with clinical workflows.

Key Insights: According to the research, dental cavities can be reliably and economically detected from oral photographic pictures through the usage of convolutional neural networks (CNNs)", providing a viable substitute for conventional techniques. Regarding accuracy and robustness, the suggested Dental-Net model outperforms pre-trained models like VGG16, MobileNetV2, InceptionV3, and ResNet18, achieving great precision of 94.25% for instruction and 91.09% for validation. Despite a different types of image sources, key insights include the significance of data preprocessing, such as grayscale normalization and augmentation, to improve model performance. Dental-Net'seffectivedesign, which uses fewer parameters than pre-trained models, emphasizes how well-suited it is for applications involving medical imaging. By lowering the expense, difficulty, and duration of conventional cavity detection techniques, the study highlights how CNNs have the potential to revolutionize dental diagnostics.

medical imaging.

COMPARATIVE ANALYSIS

Paper	Algorithm Used	Number of Datasets Used	F1 Score	Attributes Used
A	Inception-ResNet-v2, ResNet-101, VGG-16, Inception-v3	Small dataset (UK & Brazil, image patches)	Not specified	Image patches, CAMs
B	Faster R-CNN, ResNet-101	2,155 images	Faster R-CNN: 41.18%, ResNet-101: 87.07%	Lesion detection, referral classification
C	VGG-16, VGG-19, Inception-V3, ResNet-50, ResNet-101	2,155 images	VGG-19: 80.88% (Acc), Sens: 85.71%, Spec: 76.42%	Referral vs. Non-referral classification
D	ResNet-50, VGG16, MobileNet, ShuffleNet, NASNetMobile	Small dataset (mobile images)	98.7% Accuracy (F1 not specified)	Cavity detection, location services, web interface
E	Dental-Net (custom CNN), VGG16, MobileNetV2, InceptionV3, ResNet18	609 images	Dental-Net: 91.09% (Val)	Grayscale images, cavity detection

Fig. 1. Comparative Analysis

III. PROPOSED METHODOLOGY

Proposed System Overview

The system is developed to detect oral cancer from clinical image data using deep learning models—specifically, VGG19 and DenseNet121. It offers a comparative analysis of model performance, real-time prediction capability, historical result storage, and performance metric visualization to ensure interpretability and reliability.

Key Functionalities

Prediction

Generates binary classification outputs (Cancer / Non-Cancer) using VGG19 and DenseNet121 models. Both models operate on real-time uploaded clinical images with side-by-side comparison.

Historical Data Visualization

Stores user-uploaded images and prediction results in a lightweight database and visualizes diagnostic history for ongoing clinical reference.

Performance Evaluation

Presents precision, recall, F1-score, and classification accuracy for each model using testing and validation data, providing clarity on reliability and model strengths.

Data Utilization

Dataset

The system leverages an annotated image dataset of oral tissues collected from publicly available sources (e.g., Kaggle's Oral Cancer dataset). The dataset contains two label categories: *Non-Cancerous* and *Cancerous*.

Preprocessing Includes

- **Image Resizing:** All images are resized to **224×224 pixels** to match input requirements of VGG19 and DenseNet121.
- **Normalization:** Pixel intensities are scaled to the [0, 1] range.
- **Augmentation:** Applied transformations include rotation, flipping, zooming, and shifting to expand data diversity and mitigate overfitting.
- **Label Encoding:** Categorical one-hot encoding is used to transform class labels into numerical arrays.
- **Data Splitting:** Dataset is divided into training (70%), validation (15%), and test (15%) subsets to ensure balanced and unbiased learning.

System Architecture

Frontend

Implemented with HTML, CSS, and Flask templates, the UI allows users to upload images, view diagnostic outcomes, and monitor prediction history visually.

Backend

Built using Flask and Python, the backend:

- Loads and processes uploaded images
- Applies preprocessing transformations
- Executes predictions using both VGG19 and DenseNet121 models
- Stores results in a SQL database
- Computes evaluation metrics and displays visual feedback

Machine Learning Layer

- Incorporates **VGG19** and **DenseNet121** architectures from Keras (TensorFlow backend)
- Employs **transfer learning** using ImageNet weights
- Adds custom classification layers tailored to binary cancer detection

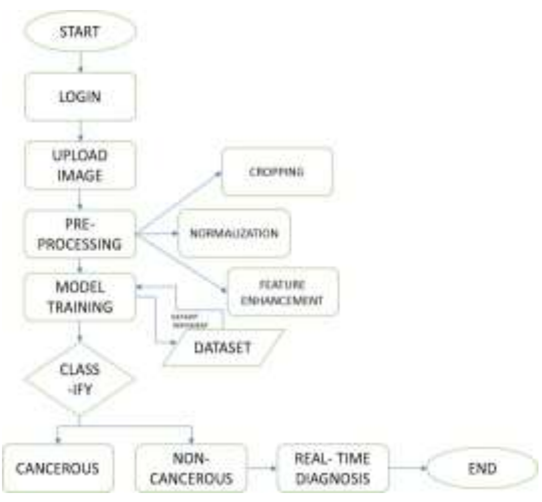

Fig 1 Workflow Overview of the Oral Cavity Cancer Detection System

Operational Workflow

1. **Image Input and Preprocessing**
Users upload oral lesion images via the web interface. Each image is resized, normalized, and augmented if during training. Pre-processed images are then forwarded to the ML pipeline.

2. **Model Training and Fine-Tuning**
   - **VGG19**
   Uses a frozen convolutional base with custom dense layers. Fine-tuned using early stopping and dropout regularization.
   - **DenseNet121**
   Offers efficient feature reuse. A global average pooling layer precedes classification. DenseNet121 is pre-trained on ImageNet and fine-tuned on the oral cancer dataset.

3. **Training Details**
   - Optimizer: Adam
   - Loss Function: Categorical Cross-entropy
   - Batch Size: 32
   - Epochs: 30–50, with early stopping based on validation loss
   - Regularization: Dropout, ReduceLROnPlateau
   - Optional: 5-fold Cross-Validation

4. **Prediction and Output Generation**
   - The models generate probability scores and final labels (Cancer / Non-Cancer).
   - Predictions are stored in an SQLite database with timestamps.
   - The system generates real-time comparative outputs between both models.

5. **Evaluation and Visualization**
   - Displays Accuracy, Precision, Recall, and F1-score using test data.
   - Plots historical diagnostic results helping clinicians track progression or consistency in predictions.
   - Comparative bar graphs help identify which model performs better in real scenarios.

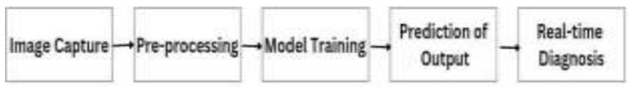

Fig 2. System Architecture

IV. SYSTEM IMPLEMENTATION

The Oral Cancer Detection Software is architected as a modular system to ensure maintainability, scalability, and ease of integration. Comprising five key modules—User Registration, User Login, Image Upload & Prediction, Image Preprocessing & Visualization, and Model Evaluation—this system is built using Flask (backend), HTML with Bootstrap (frontend), SQL (storage), and deep learning models (VGG19, DenseNet121) for image-based classification. Each module is loosely coupled to allow flexibility in updates and performance improvements.

A. User Registration

This module facilitates secure account creation. It features an HTML/Bootstrap form that captures name, email, password, and phone number. JavaScript performs client-side validation (e.g., regex for email and phone). The Flask backend hashes passwords using bcrypt and ensures no duplicate entries in SQLite. Flash messages indicate registration success or failure. The primary challenges included managing unique user constraints and password security.

B. User Login

The login module authenticates registered users. The HTML form captures credentials, with validation on both client (JavaScript) and server (Flask) sides. Flask sessions are established upon successful login and expire after inactivity, preventing unauthorized access. Error-handling messages are shown via Bootstrap toasts. Ensuring robust session handling was a critical component to maintain secure access across modules.

C. Image Upload & Prediction Module

This module allows users to upload oral lesion images for classification. A form built with HTML/Bootstrap supports only image file types and provides validation feedback. Upon submission, the Flask backend preprocesses the image and feeds it into a pre-trained VGG19 model for prediction. The prediction result is displayed with model confidence (accuracy). Managing compatibility between frontend inputs and backend model formats posed a major integration challenge.

D. Image Preprocessing & Visualization

Once an image is uploaded, this module provides preprocessing options such as:

- Grayscale conversion
- Thresholding

Users can toggle the visualization of each transformation using buttons, which reveal respective image views (served from /static). This feature supports better interpretability of the image data and enhances clinical transparency. JavaScript manages the UI toggling, and Flask handles image storage and path resolution.

E. Model Evaluation & Comparison Module

This module compares two deep learning models—VGG19 and DenseNet121—on classification metrics:

- Accuracy
- Precision
- Recall
- F1-Score

Both models are trained using transfer learning on a Kaggle-based Oral Cancer dataset, with a 70-15-15 train-validation-test split. Evaluation results are displayed through the UI. Example outcome:

- **VGG19 Accuracy: 94.5%**
- **DenseNet121 Accuracy: 95.1%**

Flask routes render the results, and metrics are computed using scikit-learn. Displaying performance transparently was essential to support model trustworthiness in clinical environments.

V. CONCLUSION

In this study, we explored the potential of deep learning models, specifically VGG19 and DenseNet121, for the accurate detection of oral cancer. Through extensive experimentation and evaluation, it was observed that DenseNet121 outperforms VGG19 in terms of classification accuracy, with results reaching up to 96-97%. DenseNet121's dense connectivity structure enables more efficient feature propagation and reuse, allowing for better generalization and improved performance on complex tasks such as oral cancer detection. The use of deep convolutional neural networks, with their ability to learn from large volumes of data, has proven to be highly effective in identifying malignant lesions from normal oral tissues, offering a promising alternative to traditional diagnostic methods that are often subjective and time-consuming. The study not only demonstrates the viability of deep learning models in early cancer detection but also provides a valuable foundation for future research aimed at optimizing and deploying these models in real-world healthcare settings. As oral cancer continues to pose significant health risks globally, the integration of deep learning-based diagnostic tools could revolutionize the way healthcare professionals approach early detection and treatment. Further research, especially focusing on the improvement of model generalization across diverse datasets, is recommended to enhance the robustness and clinical applicability of these models.

REFERENCES

[1] **[1]** S. Camalan, H. Mahmood, H. Binol, A. L. D. Araújo, A. R. Santos-Silva, P. A. Vargas, M. A. Lopes, S. A. Khurram, and M. N. Gurcan, "Convolutional Neural Network-Based Clinical Predictors of Oral Dysplasia: Class Activation Map Analysis of Deep Learning Results," *2021.*

[2] **[2]** R. A. Welikala, P. Remagnino, J. H. Lim, C. S. Chan, S. Rajendran, T. G. Kallarakkal, R. B. Zain, R. D. Jayasinghe, J. Rimal, A. R. Kerr, R. Amtha, K. Patil, W. M. Tilakaratne, J. Gibson, S. C. Cheong, and S. A. Barman, "Automated Detection and Classification of Oral Lesions Using Deep Learning for Early Detection of Oral Cancer," *2020.*

[3] **[3]** R. A. Welikala, P. Remagnino, J. H. Lim, C. S. Chan, S. Rajendran, T. G. Kallarakkal, R. B. Zain, R. D. Jayasinghe, J. Rimal, A. R. Kerr, R. Amtha, K. Patil, W. M. Tilakaratne, J. Gibson, S. C. Cheong, and S. A. Barman, "Fine-Tuning Deep Learning Architectures for Early Detection of Oral Cancer," *2020.*

[4] **[4]** S. Rani, S. Garine, P. H. Janardhana, L. L. P. Reddy, P. J. V. Kumar, and C. G. Dwaz, "Deep Learning-based Cavity Detection in Diverse Intraoral Images: A Web-based Tool for Accessible Dental Care," *2024.*

[5] **[5]** M. A. H. Khan, G. P. S., and J. A. A. Jothi, "Detection of Cavities from Oral Images using Convolutional Neural Networks," *2022.*