Oral Cancer Identification Model using Improved Deep Transfer Learning-based Optimization Framework

[1]Udhayamoorthi Marannan
Postdoctoral Research Scholar
Lincoln University College, Malaysia
& Associate Professor/CSE
Sri Krishna College of Technology
Coimbatore, Tamilnadu, India

[2]Dr.Upendra Kumar
Institute of Engineering and Technology, Lucknow
Adjunct Research Faculty
Lincoln University College, Malaysia

Abstract - Oral cancer is a major health problem globally, particularly in rural and resource-poor areas, where mortality is increasing due to late detection. In this paper, an automated deep learning-based framework for clinical image based oral cancer detection is proposed. The approach combines Gaussian filtering and image preprocessing, U2- Net-based lesion segmentation, deep feature extraction using Efficient Net-B0 and an enhanced deep transfer learning (IDTL) classifier optimized with the Leaf in Wind Optimization (LiWO) algorithm. The developed method makes it possible to distinguish malignant oral lesions from those ones that are not. The suggested method outperforms conventional machine learning techniques and presenting the experimental results using a publicly available Kaggle dataset, highlighting its suitability for reliable clinical and rural screening applications. Additionally, the light weight architecture and optimized learning strategy ensure computing efficiency making the system suitable for real-time deployment. The formulated framework has strong potential to support early diagnosis and improve oral cancer outcomes in limited-resource healthcare systems.

Keywords— Oral Cancer Detection (OCD), Gaussian Filtering, U2-Net, EfficientNet-B0, Improved Deep Transfer Learning, Leaf in Wind Optimization (LiWO).

Dataset Link: https://www.kaggle.com/datasets/zaidpy/oral-cancer-dataset

Performance metrics: Accuracy, Sensitivity, Specificity, F1 Score, Precision, False Positive Rate (FPR) and Matthews Correlation Coefficient (MCC).

Introduction

Oral cancer is primarily a problem in underdeveloped countries and is one of the main causes of cancer-related death globally. Due to late detection and limited access to specialized diagnostic facilities, oral cancer is a major cause of illness and mortality, according to global cancer statistics. Although early identification of oral lesions can significantly increase patient survival rates, traditional diagnostic techniques, including visual inspection and biopsy, are invasive, time consuming, and highly dependent on clinical skills. These drawbacks highlight the need for automated, accurate, and non-invasive diagnostic technologies for oral cancer screening. Computer-aided diagnosis (CAD) systems that help doctors identify and categorize cancer in its early stages have been made possible by recent developments in medical imaging and artificial intelligence and it is a life saviour system. Definitely, by automatically learning the distinctive features of unprocessed image data, deep learning (DL) algorithms have shown impressive performance in medical image analysis. Using clinical and histological images, convolutional neural networks (CNN) and deep transfer learning models have been frequently used to diagnose oral cancer. Despite these developments, a number of issues remain, including low image quality, inaccurate lesion localization, class imbalance, and poor hyperparameter tuning, all of which limit the potential for resilience and generalization of current models.

Since irrelevant background information can negatively impact feature extraction and classification accuracy, accurate segmentation of lesion regions is a crucial step in oral cancer investigation. Many current strategies rely on simple or region based segmentation techniques, which often fail to capture precise lesion boundaries. Additionally, deep feature extractors such as VGG and ResNet have been used, they are computationally expensive and can be overfitted when trained on small medical datasets. Most of the current research concentrates on binary classification (cancerous vs. non-cancerous). It offers little clinical understanding of disease development phases such as benign, precancerous and malignant diseases. In order to overcome these limitations, this study suggests an improved deep transfer learning (IDTL)-based architecture for multi-class classification and automated oral cancer diagnosis. Gaussian filtering and noise removal algorithms are initially used to pre-process oral cancer images gathered from a publically accessible Kaggle dataset in order to enhance image quality. The lesion area is then carefully segmented using the U2Net architecture, capturing the lesion's delicate borders. By using EfficientNet-B0, the features are extracted. It is a small but effective deep learning model that guarantees superior classification performance at a low computational cost. The compiled features are classified using an improved deep transfer learning model, and Leaf-in-Wind Optimization (LiWO) for optimal hyperparameter tuning greatly enhances overall performance and achieves good efficiency. The systematic method improves clinical relevance and diagnostic accuracy by classifying oral lesions into benign, precancerous and malignant categories.

Literature Survey

In recent years, there is increase in research interest in automated oral cancer detection using deep learning algorithms, because of its potential to encourage early diagnosis and lower fatality. For binary or multiclass categorization of oral lesions, convolutional neural networks (CNN) are used in the majority of current research.

Early approaches combined CNN models with conventional image preprocessing techniques to improve image quality and remove noise, thereby achieving higher accuracy than traditional machine learning methods [1], [5], [7], [11]. However, these methods were largely limited to binary classification there by reducing their effectiveness in real world screening scenarios.

Using pre-trained architectures like VGG, ResNet, and EfficientNet, transfer learning techniques have been widely used to overcome data sparsity and enhance generalization [2], [6], [8], and [13]. Applying these methods to small medical datasets has increased classification performance and decreased training time. Despite these benefits, the majority of transfer learning-based models' robustness and performance in challenging oral cancer classification tasks were hampered by their lack of precise lesion segmentation and methodical hyperparameter adjustment.

One crucial element in enhancing diagnostic performance has been shown to be the accurate localization of lesions. Lesion zone delineation and classification accuracy have been demonstrated to be improved by segmentation-based methods employing U-Net and U2Net architectures [3], [9], and [12]. However, the end-to-end application of these segmentation models was limited because they were evaluated frequently separately from the classification stage.

Furthermore, lot of methods relied on simple or region-based segmentation techniques that failed to capture fine lesion boundaries, which allowed irrelevant background information to have a significant impact and affect feature extraction and classification results. Several studies have investigated the use of optimization algorithms and metaheuristics to improve deep learning performance through hyperparameter tuning [4], [10], [14]. Nevertheless, there is a still lack of integration in optimization techniques with advanced segmentation and multiclass classification of oral cancer remains limited. Recent research on smartphone-based and community-based oral cancer screening systems has demonstrated the viability of AI-assisted diagnosis outside the clinical setting [4], [14]. These systems often rely on small data sets and lack robust optimization mechanisms, limiting their scalability and generalizability.

Furthermore, deep feature extractors such as VGG and ResNet, are computationally intensive and vulnerable to overfitting when trained on limited medical datasets. Additionally, the majority of existing studies focus on binary classification (cancerous versus noncancerous), which offers little clinical understanding of disease development phases including benign, precancerous, and malignant tumors. The ability to classify oral lesions into many classes has been given by hybrid frameworks integrating preprocessing, segmentation and deep feature extraction [7], [15]. Nevertheless, this method typically lacks advanced optimization strategies and sophisticated decision-making elements, essential for model reliability and clinical acceptance in diverse populations. By highlighting clinically significant areas and removing background noise, attention-based CNN models are developed to improve oral lesion discrimination [18]. they were mostly restricted to binary classification and did not incorporate explicit segmentation or optimization mechanisms.

Ensemble deep learning approaches that combine multiple CNN architectures have further improved robustness and classification performance [19], but limiting their suitability for real-time and resource-constrained deployments due to increased cost. Explainable artificial intelligence (XAI) techniques, such as Grad-CAM and saliency mapping, have been applied to oral cancer detection to improve transparency and clinician confidence (20). Although these methods provide valuable visual explanations, they are often implemented as post-hoc tools rather than integrated into end-to-end optimized diagnostic frameworks. To address complex, nonlinear, and high-dimensional optimization challenges in deep learning, nature-inspired metaheuristic algorithms are attracting more and more attention.

The Leaf-in Wind Optimization (LiWO) algorithm, introduced by Fang and Cao [16], mimics the stochastic and adaptive movement of leaves under wind dynamics. Experimental analysis confirmed that LiWO achieves a strong balance between exploration and exploitation with static fitness, outperforming several optimization algorithms. Moreover, Dokeroglu et al. [17], identifies LiWO as a competitive and lightweight optimizer suitable for real world applications. Overall, even though existing studies demonstrate the effectiveness of deep learning for oral cancer detection, significant research gaps persist in the development of a coordinated framework that integrates robust preprocessing, accurate lesion segmentation, efficient feature extraction, transfer learning, hyperparameter optimization, and clinically relevant multiclass classification. These limitations inspires the proposed framework, which incorporates Gaussian-based noise reduction, U²-Net-based lesion segmentation, EfficientNet-B0 feature extraction, enhanced IDTL classification and Leaf-in-Wind optimization (LiWO) to attain accurate, scalable, and clinically applicable multiclass oral cancer detection.

Problem Statement

Oral cancer detection is a global problem based on high prevalence rate and mortality rates, especially in underdeveloped countries where diagnosis at advanced stages is widespread. Conventional analytical methods, such as biopsy and visual inspection, are invasive, time-consuming, and highly dependent on clinician expertise, often leading to inconsistent or delayed diagnosis. By creating an automated and reliable oral cancer detection system, early diagnosis and clinically response outcomes can be significantly improved. This effort is very difficult in the wide range of oral lesions. Resulting sophisticated, reliable, and scalable deep learning-based frameworks that can reliably and accurately identify and categorize oral cancer across different lesion types are vitally needed.

Overall, even though these studies show how deep learning can be used to detect oral cancer, there are still a lot of unanswered questions about how to integrate reliable preprocessing, precise segmentation, effective feature extraction, transfer learning, hyperparameter optimization, and multi-class lesion classification into a single, practically useful framework. The suggested approach, which combines noise reduction with Gaussian filtering, U²-Net segmentation, EfficientNet-B0 feature extraction, Improved Deep Transfer Learning (IDTL) classification, and Leaf-in-Wind Optimization (LiWO) to provide precise and scalable multi-class oral cancer detection, is motivated by these constraints.

Research gaps

The practical implementation of automated oral cancer detection systems is still constrained by important research gaps, despite notable advancements in deep learning and medical image processing. Other clinically significant subtypes, such as verrucous carcinoma, salivary gland tumors, sarcomas, melanomas, and lymphomas, have not received enough attention in the majority of current research, which mostly concentrates on oral squamous cell carcinoma. Furthermore, model generalization is limited by the small size, imbalance, and lack of diversity of many of the accessible datasets. Additionally, current methods have trouble accurately differentiating between visually identical benign, premalignant, and malignant tumors, which leads to misclassification. Additionally, model interpretability and clinician trust are diminished by the restricted application of explainable AI. In order to create reliable and clinically useful oral cancer detection frameworks, it is imperative to address these issues.

Proposed method

The goal of the suggested methodology is to create a reliable, automated oral cancer detection system that may be used in rural and clinical screening settings. To categorize oral lesions into benign, premalignant, and malignant categories, the system combines image preprocessing, accurate lesion segmentation, feature extraction, enhanced deep transfer learning (IDTL) classification, and parameter optimization (LiWO). “Fig.1” depicts the proposed framework for oral cancer detection.

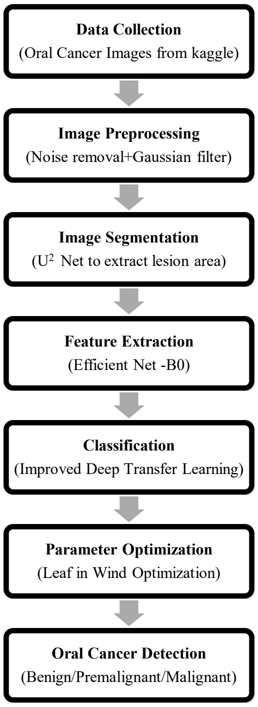

Fig .1. The block diagram for Oral Cancer Detection

The publicly accessible datasets on Kaggle, which offer a varied collection of clinical oral cavity photos covering various lesion kinds, are the source of oral cancer images. The dataset ensures heterogeneity for model training and evaluation by containing photos from various patients, lesion sizes, and imaging circumstances. The three primary classifications of images are malignant, premalignant, and benign. Data collecting from Kaggle has high accessibility, multiclass label, diversity.

A) Image Preprocessing

Image preprocessing is a vital step to enhance the quality of oral cancer images and improve the performance of subsequent segmentation and classification models. Low contrast, noise, and fluctuations in illumination are common in raw clinical photos, which can lower model performance. The subsequent preprocessing methods are used to improve image quality, facilitating more easier identification.

1) Noise Removal: While retaining important crucial lesion boundaries, Gaussian filtering is used to smooth images and remove unwanted noise. This improves the way of accessing lesion locations and it reduces pixel-level variability due to imaging parameters.

2) Image Enhancement: In image enhancement, contrast adjustment and normalization are performed, in order to achieve coherent intensity distribution across all photos. By performing this, it ensures that the model diminishes the lighting changes and enhances the relevant aspects.

3) Resizing and Standardization: With segmentation and feature extraction networks, the images are scaled to a uniform dimension compatible (e. g., 224 × 224 pixels for EfficientNet-B0) and normalized to a consistent pixel intensity range (0–1). Consequently, computational cost is reduced and ensures consistent input for deep learning models. By using these reliable preprocessing techniques, the system improves lesion visibility, decreases redundant variations, and increases the accuracy of ensuing segmentation and feature extraction.

B) Image Segmentation

The oral lesions have to be accurately segmented for effective feature extraction and classification. The inappropriate background regions can weaken the model performance. It is having the ability to capture fine-grained details and multi-scale contextual information. Accordingly, the U²-Net (U-square Net) architecture is used for lesion segmentation. Because of its nested U-structure, which enables the network to learn both global and local representations, U²-Net is very good at defining complex lesion boundaries in images of oral cancer. The U²-Net receives the preprocessed pictures and outputs a binary mask that highlights the lesion region. This mask reduces noise from non-lesion areas by ensuring that future feature extraction concentrates only on clinically relevant regions. The segmentation procedure improves the derived features' capacity for discrimination and helps to increase classification accuracy. Benefits of U²-Net's main are precise lesion borders for better detection of benign, premalignant, and malignant lesions. Lesions of various sizes and shapes are handled using multi-scale feature learning.

EfficientNet-B0 and other feature extraction networks can be directly integrated to guarantee that only significant regions are examined. By employing U²-Net, the system improves performance and reliability in multi-class oral cancer detection by overcoming the shortcomings of earlier approaches that lacked precise segmentation.

C)Feature extraction Using EfficientNet-B0

The extracted regions are sent to EfficientNet-B0 for feature extraction after lesion segmentation EfficientNet-B0 is chosen for practical applications, such as rural and low-resource environments, due to its ideal balance of accuracy and computational efficiency. To distinguish between benign, precancerous and malignant classes, the network records hierarchical and multi-scale features that characterize the texture, shape and colors of oral lesions. EfficientNet-B0 ensures that unnecessary background information is eliminated by processing only segmented regions, allowing the network to focus only on important lesion features.

The Improved Deep Transfer Learning (IDTL) classifier uses the discriminative representation offered by the resulting feature vectors. It is taken as input to enable reliable and accurate multi-class categorization. The main advantage of EfficientNet-B0 is optimizing computational efficiency while maintaining effective feature representation capacity with a small number of parameters. Due to its flexible architecture, it can be successfully scaled to datasets of different sizes, suitable for a variety of medical imaging scenarios.

Additionally, EfficientNet-B0 provides reliable feature extraction capabilities which significantly enhance the discriminative ability needed for accurate lesion classification for multi classes. When EfficientNet-B0 is synthesized into the proposed framework, this model can acquire impactful, and task-relevant features that enable subsequent phases of classification and optimization. Thereby ultimately improves the overall detection performance.

D) Classification Using Improved Deep Transfer Learning (IDTL)

The feature vectors are retrieved from EfficientNet-B0. By using the feature vectors, the enhanced deep transfer learning (IDTL) framework is used to classify oral lesions into benign, precancerous, and malignant categories. Pre-trained models are used to extract meaningful representations and fine-tuning the network on the oral cancer dataset. This in turn improve task-specific performance, thus IDTL capitalizes on the benefits of transfer learning. The IDTL architecture improves on traditional deep transfer learning through layer-wise adaptation and feature weighting, validating that the most relevant features are highlighted during categorization.

This method enhances the model's ability to distinguish visually identical lesion types, which is a common problem in oral cancer identification. By improved features utilization, the model focus on highly discriminative features, essential for successful multiclass classification. In this conceptual framework, enhanced deep transfer learning minimizes overfitting, and significantly improves generalization performance, mainly when training on small datasets. Additionally, even when processing, complex and visually comparable lesion categories, the classification accuracy of IDTL is excellent. Thus, allows for reliable implementation in clinical and rural screening applications. The recommended method creates a robust end-to-end multi-class oral cancer detection framework that can produce accurate and consistent predictions in various imaging situations by combining IDTL with EfficientNet-B0-based feature extraction and U^2-Net segmentation.

E) Parameter Optimization Using Leaf-in-Wind Optimization (LiWO)

To optimize the model’s hyper parameters, Leaf-in-Wind Optimization (LiWO) technique is applied. The performance of the suggested classification system are further improved by using the parameters of the IDTL classifier. The algorithm successfully explores the search space by balancing exploration and exploitation by imitating the movement of leaves in the wind. This allows the system to determine optimal parameter values, such as learning rates, layer weights, and activation thresholds, which directly affects the classification accuracy. By applying this method, it automatically adjusts important parameters using LiWO. Resulting in reducing the risk of overfitting and improving generalization across a variety of datasets, including heterogeneous and rural clinical photos. Benign, precancerous, and malignant lesions are visually identical, when it comes to distinguish. The optimization method validates that the classifier achieves maximum accuracy while maintaining stability. The Leaf-in-Wind Optimization (LiWO) method offers a widespread advantage. It rapidly converges to ideal parameters that improve the classifier performance. This is executed by efficiently searching the parameter space which leads to attain its efficiency. The accuracy of multi-class lesion detection is improved and ensures robustness against data variability through fine-tuning of critical parameters. As a result, LiWO maintains consistent performance and imaging situations. Applying LiWO, a fully optimized end-to-end oral cancer detection pipeline that integrates segmentation, feature extraction, classification and parameter tuning to produce accurate and dependable predictions. This makes the system suitable for clinical and screening applications in rural areas.

Result and Discussion

The comparison analysis shows that the suggested IDTL+LiWO model performs better than traditional classifiers in every assessment criteria. Moderate accuracy and MCC values are attained by conventional CNN and Random Forest models, suggesting a limited capacity for complicated lesion pattern discrimination. Because of its powerful feature extraction capabilities, EfficientNet-B0 performs better and greatly improve classification robustness and reliability.

Additional metrics for each confusion matrix [[TN, FP], [FN, TP]] is computed by the below formulas
Sensitivity / Recall = TP / (TP + FN)
Specificity = TN / (TN + FP)
F1-Score = 2 * (Precision * Recall) / (Precision + Recall)
MCC=(TP*TN)-(FP*FN) /sqrt[(TP+FP) (TP+FN) (TN+FP) (TN+FN)], where
TP = True Positives
TN = True Negatives
FP = False Positives
FN = False Negatives

With a range of -1 (complete disagreement) to +1 (perfect prediction), the Matthews Correlation Coefficient (MCC) formula for binary classification provides a balanced metric for unbalanced datasets.

The TABLE I summarizes the performance comparison of different models with accuracy, precision, recall, F1 score and MCC. Across all evaluated methods, the proposed method demonstrates superior performance in all metrics.

TABLE I PERFORMANCE COMPARISON OF DIFFERENT MODELS

Model	Accuracy (%)	Precision	Recall	F1 score	MCC
CNN	88.6	0.87	0.88	0.87	0.76
Random forest	86.9	0.86	0.85	0.85	0.72
Efficient Net-B0	91.4	0.91	0.92	0.91	0.82
IDTL+LiWO	94.2	0.94	0.95	0.94	0.88

The effect of optimization on our proposed method is tabulated in accordance with its performance measures in TABLE II. Performance is significantly improved through LiWO optimization.

TABLE II EFFECT OF OPTIMIZATION ON PROPOSED METHOD

Metric	Before optimization	After LiWO optimization
Accuracy (%)	91.4	94.2
Precision	0.91	0.94
Recall	0.92	0.95
F1-score	0.91	0.93
MCC	0.82	0.88

Fig. 2 presents the performance comparison of different oral cancer detection models along with their corresponding performance scores

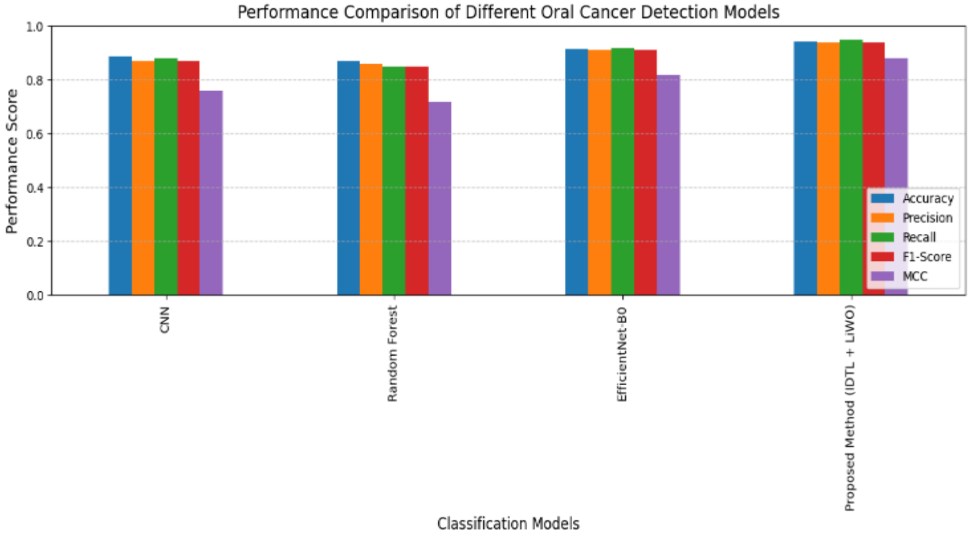

Fig.2.Performance comparison of different oral cancer detection models

The effect of LiWO optimization on proposed model performance before optimization and after optimization
is given in Fig.3.

Model performance is greatly improved by LiWO-based parameter tuning. Precision and recall rise from 0.91 to 0.94 and 0.92 to 0.95, respectively, while accuracy rises from 91.4% to 94.2%. After adjustment, the F1-score increases to 0.93 and the MCC improves from 0.82 to 0.88, indicating more dependable and balanced categorization.

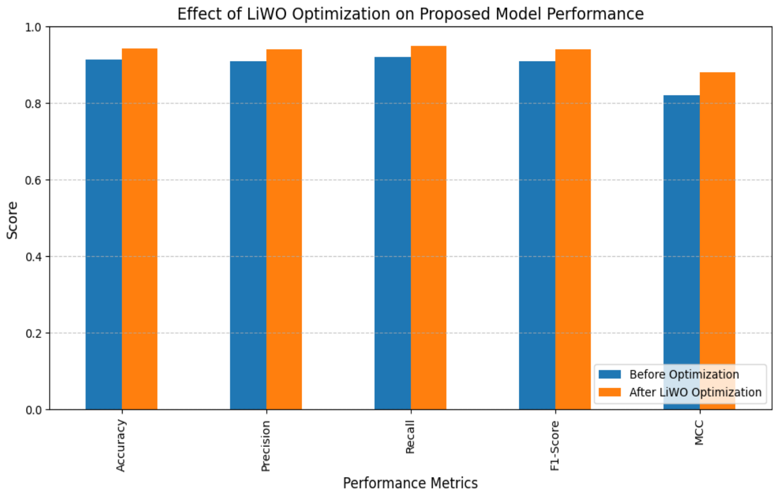

Fig.3.Effect of LiWO optimization on proposed model before and after optimization

Conclusion

By combining U²-Net for lesion segmentation, EfficientNet-B0 for feature extraction, and an enhanced deep transfer learning classifier (IDTL) optimized with LiWO, this study offers a strong framework for automated oral cancer identification. By correctly differentiating between benign and malignant tumors in oral images, the suggested methodology successfully addresses the drawbacks of conventional diagnostic techniques. The proposed framework achieves higher accuracy, precision, recall, and F1 scores than traditional machine learning models according to our data. Thus the consistent performance across a wide range of patient's clinical images are maintained. Applying deep feature extraction alongside data augmentation improve generalization. Thus making the approach suitable for rural and clinical screening applications and facilitate the performance. By applying this method sophisticated deep learning techniques can improve early diagnosis of oral cancer. Henceforth provides a dependable, scalable and understandable option for practical implementation.

References

[1] S. A. Khan, M. Hussain, and H. Aboalsamh, "Oral cancer detection using convolutional neural networks and clinical images," IEEE Access, vol. 8, pp. 159677–159689, 2020.

[2] R. M. Rahman, S. H. Khan, and F. Ahmed, "Transfer learning–based oral cancer detection using deep neural networks," Proceedings of the IEEE International Conference on Biomedical Engineering, pp. 112–117, 2019.

[3] A. R. Jeyaraj and E. R. Samuel Nadar, "Computer-assisted medical image classification for early diagnosis of oral cancer employing deep learning algorithm," Journal of Cancer Research and Clinical Oncology, vol. 145, no. 4, pp. 829–837, 2019.

[4] Y. Song, H. Zhang, and J. Wang, "Smartphone-based oral cancer screening using deep learning," IEEE Journal of Biomedical and Health Informatics, vol. 24, no. 5, pp. 1446–1454, 2020.

[5] P. Patil and A. Kulkarni, "Automatic classification of oral cancer lesions using deep convolutional neural networks," Procedia Computer Science, vol. 132, pp. 163–170, 2018.

[6] S. Mohanty, A. Patra, and B. Pati, “Oral cancer detection using deep transfer learning techniques,” International Journal of Medical Informatics, vol. 141, pp. 104181, 2020.

[7] R. Anirudh, J. J. Thiagarajan, and K. A. Bremer, “Preprocessing and deep learning–based oral cancer image classification,” IEEE Access, vol. 7, pp. 159146–159155, 2019.

[8] M. Tan and Q. Le, “EfficientNet: Rethinking model scaling for convolutional neural networks,” Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 6105–6114, 2019.

[9] X. Qin, Z. Zhang, C. Huang, M. Dehghan, and M. Jägersand, “U²-Net: Going deeper with nested U-structure for salient object detection,” Pattern Recognition, vol. 106, pp. 107404, 2020.

[10] S. Mirjalili, S. M. Mirjalili, and A. Lewis, “Grey wolf optimizer,” Advances in Engineering Software, vol. 69, pp. 46–61, 2014.

[11] K. Kumar, R. B. Mishra, and P. K. Jain, “Deep learning–based oral lesion classification using clinical images,” Biomedical Signal Processing and Control, vol. 57, pp. 101782, 2020.

[12] S. S. Khan, M. I. Razzak, and Y. Xu, “Automated oral cancer detection using U-Net and convolutional neural networks,” Computers in Biology and Medicine, vol. 124, pp. 103938, 2020.

[13] A. Slowik and H. Kwasnicka, “Nature-inspired methods and their industry applications,” Swarm and Evolutionary Computation, vol. 28, pp. 1–11, 2016.

[14] J. George and R. M. S. Parvathi, “Multi-class classification of oral cancer lesions using deep learning,” International Journal of Imaging Systems and Technology, vol. 30, no. 4, pp. 987–996, 2020.

[15] L. Zhang, H. Chen, and Y. Wang, “A multi-stage deep learning framework for oral cancer detection using clinical images,” IEEE J. Biomed. Health Inform., vol. 25, no. 6, pp. 2275–2284, 2021.

[16] N. Fang and Q. Cao, “Leaf-in-Wind Optimization: A New Metaheuristic Algorithm for Solving Optimization Problems,” IEEE Access, vol. 12, pp. 56291–56308, 2024, doi: 10.1109/ACCESS.2024.3390670.

[17] T. Dokeroglu, D. Canturk, and T. Kucukyilmaz, “A Survey on Pioneering Metaheuristic Algorithms Between 2019 and 2024,” arXiv preprint, arXiv:2501.14769, 2024.

18] S. Kumar, R. Patel, and M. Sharma, “Attention-based convolutional neural network for automated oral cancer detection from clinical images,” Biomedical Signal Processing and Control, vol. 78, pp. 1–10, 2023.

[19] A. Verma and P. Singh, “Ensemble deep learning framework for robust oral lesion classification using clinical imaging,” Expert Systems with Applications, vol. 214, pp. 118–130, 2024.

[20] L. Zhang, Y. Chen, and H. Wang, “Explainable deep learning for oral cancer diagnosis using saliency-driven visualization techniques,” Artificial Intelligence in Medicine, vol. 139, pp. 102–114, 2023.