RESEARCH ARTICLE

Editorial Process: Submission:05/26/2025 Acceptance:05/12/2026 Published:05/18/2026

Oral Cancer Diagnosis Using an Optimized InceptionV3 Model Powered by the Aquila Metaheuristic Algorithm

Sabura Banu Urundai Meeran[1]*, Kavitha S.[1], Chinthamani B.[2]

Abstract

Objective: Early and accurate detection of oral cancer plays a pivotal role in improving patient prognosis and survival rates. Deep learning (DL) models have shown promise in automating medical image classification; however, performance optimization remains a challenge due to complex network configurations and hyperparameter dependencies. This study introduces an enhanced diagnostic framework combining the InceptionV3 convolutional neural network with the Aquila Optimizer (AO), a nature-inspired metaheuristic algorithm, to achieve superior classification accuracy in identifying oral cancer lesions. **Methods:** A standardized dataset of labeled oral lesion images, including both benign and malignant cases captured via mobile and intraoral cameras, was used for training. The InceptionV3 model, initially pre-trained, was fine-tuned for binary classification tasks. AO was employed to optimize the hyperparameters by defining a search space and iteratively improving model performance through accuracy maximization and loss minimization strategies. The optimized model was compared against leading architectures such as AlexNet, MobileNet, Xception, ResNet-50, and the original InceptionV3, using comprehensive performance indicators like accuracy, precision, recall, F1-score, AUC-ROC, specificity, log loss, and Matthews Correlation Coefficient (MCC). **Result:** The proposed AO-InceptionV3 model consistently outperformed the other DL architectures across all metrics. It achieved a classification accuracy of 97.80%, precision of 97.81%, recall of 97.79%, and an MCC of 0.956, while maintaining a low log loss of 0.0735 and an AUC-ROC of 99.81%. Visual analyses, including ROC curves and 3D plots, reinforced the robustness and reliability of the model in distinguishing between benign and malignant lesions with minimal inference time. **Conclusion:** The integration of the Aquila Optimizer into the InceptionV3 architecture significantly improves the diagnostic performance of DL models for oral cancer detection. The proposed framework demonstrates excellent potential for real-time clinical deployment, offering high accuracy, efficiency, and reliability, and sets a benchmark for future AI-driven cancer diagnostic systems.

Keywords: Oral Cancer Detection- Deep Learning- Inception V3- Aquila Optimizer- Medical Image Classification

Asian Pac J Cancer Prev, **27 (5)**, 1623-1632

Introduction

Oral cancer is a major global health concern, with oral squamous cell carcinoma (OSCC) accounting for more than 90% of reported cases. Recent estimates indicate that over 370,000 new cases of oral cancer are diagnosed annually worldwide, with a disproportionately high burden in low- and middle-income countries [1]. Risk factors such as tobacco use, alcohol consumption, betel quid chewing, and poor oral hygiene significantly contribute to disease prevalence, particularly in South and Southeast Asia. Despite advances in treatment, the five-year survival rate remains low when diagnosis occurs at advanced stages, underscoring the critical importance of early and accurate detection [2].

Conventional diagnostic approaches, including visual inspection, biopsy, and histopathological analysis, are considered gold standards but suffer from several limitations. These methods are invasive, time-consuming, resource-intensive, and prone to inter-observer variability, particularly in early-stage lesions where visual cues are subtle [3, 4]. Limited access to experienced oral pathologists further exacerbates diagnostic delays, especially in rural and resource-constrained settings. As a result, there is a growing de-mand for automated, objective, and reliable computer-aided diagnostic systems to support clinicians in oral cancer screening and decision-making [5].

Recent advancements in artificial intelligence (AI) and deep learning (DL) have significantly transformed medical image analysis. Convolutional neural networks (CNNs) have demonstrated strong capabilities in learning hierarchical and discriminative features from complex image data, enabling effective classifica-tion of benign

*Department of Electrical and Electronics Engineering, Saveetha Engineering College, Saveetha Nagar, Thandalam, Chennai : 602105, Tamilnadu, India. [2]Department of Electronics and Communciation Engineering, Easwari Engineering College, Ramapuram, Chennai, Tamilnadu, India. *For Correspondence: saburasec@gmail.com*

and malignant lesions [6–8]. Several studies have explored DL-based approaches for oral cancer detection using histopathological images, clinical photographs, optical coherence tomography, and radiological scans, reporting encouraging results [9–12]. Architectures such as AlexNet, ResNet, Mo-bileNet, and Xception have been widely adopted due to their robust feature extraction abilities and adaptability to transfer learning paradigms [13–15]. Among these architectures, Inception V3 has gained particular attention for its multi-scale feature extraction capability, efficient factorized convolutions, and reduced computational complexity. Prior studies have successfully employed Inception V3 for various cancer detection tasks, including skin cancer and oral lesion classification, achieving high diagnostic accu-racy [16–18]. However, the performance of deep CNN models is highly sensitive to hyperparameter selec-tion, such as learning rate, batch size, dropout ratio, and optimizer configuration. Manual hyperparameter tuning is often heuristic, computationally expensive, and prone to suboptimal convergence, limiting mod-el generalizability and robustness [19].

To address these challenges, metaheuristic optimization algorithms have been increasingly integrated with DL frameworks to automate hyperparameter tuning and improve classification performance. Tech-niques such as Genetic Algorithms (GA), Particle Swarm Optimization (PSO), and Whale Optimization Al-gorithm (WOA) have shown promise in enhancing CNN-based medical image classifiers [20–22]. These methods improve convergence behavior and help escape local minima in highly non-convex optimization landscapes. Nevertheless, many existing approaches still suffer from premature convergence, limited ex-ploration ability, or increased computational overhead.

The Aquila Optimizer (AO) is a recently proposed nature-inspired metaheuristic algorithm modeled on the hunting strategies of Aquila (eagle) species. AO dynamically balances global exploration and local exploi-tation through four adaptive behaviors, enabling efficient search of complex optimization spaces [23]. While AO has demonstrated superior performance in engineering optimization and parameter estimation tasks [24, 25], its application in medical image analysis particularly for oral cancer detection remains largely unexplored.

Motivated by these observations, this study proposes an AO-optimized Inception V3 framework for auto-mated oral cancer diagnosis using medical images. The proposed approach integrates the powerful fea-ture extraction capability of Inception V3 with AO-based hyperparameter optimization to enhance classi-fication accuracy, robustness, and generalization. The model is evaluated against several state-of-the-art DL architectures using comprehensive performance metrics, including accuracy, precision, recall, F1-score, AUC-ROC, log loss, and Matthews Correlation Coefficient (MCC), which are particularly relevant for imbal-anced medical datasets [26]. Additionally, an ablation study is conducted to explicitly isolate and quantify the impact of AO on model performance. By leveraging optimization-driven deep learning, this work aims to provide a reliable and deployable solution for early oral cancer detection, particularly in low-resource clinical environments [27].

The remainder of this article is organized as follows. Section 2 describes the materials and methods em-ployed in this study, including the dataset characteristics, data preprocessing steps, the Inception V3 ar-chitecture, and the proposed integration of the Aquila Optimizer for hyperparameter optimization, along with model training and evaluation metrics. Section 3 presents the experimental results and comparative analysis, incorporating confusion matrix analysis, performance metric evaluation, visualization-based as-sessments, theoretical justification of the Aquila Optimizer, and an ablation study to isolate its impact. Section 4 discusses the limitations of the proposed approach and examines its feasibility for deployment in low-resource clinical environments. Section 5 provides a comprehensive discussion of the findings, highlighting the clinical relevance and advantages of the AO-optimized Inception V3 framework. Finally, the article concludes with acknowledgements, declarations, and references supporting the study.

Materials and Methods

Dataset Description

The foundation of any DL model's success lies in the effective and diversity of the dataset. For this study, a publicly available and medically validated dataset of oral cancer images was used. The dataset comprises color images of oral lesions acquired through both mobile and intraoral cameras. These images are in-tended for use in the detection of potential oral malignancies through advanced image analysis tech-niques[26]. The data collection process was carried out in collaboration with medical professionals from various hospitals and academic institutions across Karnataka, India, ensuring clinical relevance and diag-nostic accuracy. The dataset comprises high-resolution images captured through histopathological imag-ing and clinical photography. It includes two primary classes: malignant and benign. Each image was care-fully labeled by domain experts to ensure ground truth accuracy. Methods of data augmentation such as horizontal and vertical flipping, rotation, brightness variation, and scaling were applied to increase the dataset size, balance class distribution, and improve the model's generalization capability. The sample of oral Cancer is shown in Figure 1.

Methodology

Data Preprocessing

The dataset was subjected to a number of preprocessing procedures to normalize the input before train-ing

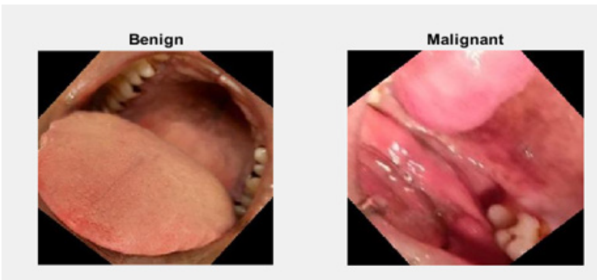

Figure 1. Sample Oral Cancer Images

and enhance feature extraction. All captured images were scaled to 299 × 299 pixels to align with the input dimensions required by the Inception V3 architecture. Normalization was applied to scale pixel val-ues to the [0,1] range. Noise removal and contrast enhancement were also performed to emphasize criti-cal diagnostic features. In order to preserve class proportion, the dataset was then separated into training (70%), validation (15%), and testing (15%) sets

Inception V3 Architecture

Inception V3 is a well-established convolutional neural network that employs factorized convolutions, auxiliary classifiers, and an efficient grid size reduction technique. It is known for balancing model com-plexity and computational efficiency, making it a strong baseline for image classification tasks. The archi-tecture integrates several convolutional and pooling layers, along with batch normalization and inception modules, facilitating multi-scale feature extraction. In this work, the pre-trained Inception V3 model, fine-tuned with the oral cancer dataset, served as the base architecture for performance benchmarking and enhancement.

Aquila Optimizer (AO) optimized Inception V3 for Oral Cancer detection

Treatment results are much improved survival rates when oral cancer is detected early and accurately. While DL models such as Inception V3 have demonstrated remarkable capabilities in image classification, their performance heavily relies on optimal hyperparameter settings. Traditional optimization methods often suffer from convergence issues or getting trapped in local minima. To address these challenges, this work introduces the integration of the Aquila Optimizer (AO) a novel nature-inspired metaheuristic algo-rithm with the Inception V3 model to improve the accuracy and robustness of oral cancer classification. The Aquila Optimizer (AO) is a nature-inspired metaheuristic algorithm modeled after the hunting strategy of eagles (Aquila species). It dynamically switches between exploration and exploitation modes based on four strategic behaviors; contour flying with brief glide, low flight with a fast assault, high soar with verti-cal stoop, and random flight with dive attack. In this study, AO was integrated with the Inception V3 model to optimize hyperparameters such as learning rate, batch size, and layer-specific weights. The op-timizer improves convergence speed and helps escape local minima, thereby enhancing classification ac-curacy and robustness. Inception V3 is a robust convolutional neural network (CNN) architecture recog-nized for its effectiveness in capturing features at multiple scales features through a combination of con-volutional blocks and inception modules. In this study, Inception V3 is fine-tuned specifically for dividing pictures of oral lesions into those that are malignant and those that are not.To enhance the learning ca-pability of the model, the Aquila Optimizer is incorporated into the training process. Inspired by the hunt-ing strategies of Aquila (eagle) species, AO intelligently balances between exploration (searching the glob-al space) and exploitation (focusing on promising regions). It uses four key strategies:

1. High soar with vertical stoop: Wide-range search to explore solutions.
2. Contour flight with short glide: Moderate search around known good solutions.
3. Low flight with quick attack: Fast local search for fine-tuning.
4. Random flight with dive attack: Enhancing the chance of escaping local minima.

The AO algorithm dynamically adjusts parameters such as learning rate, dropout rates, and weight initiali-zation during model training. This results in faster convergence, better generalization, and improved clas-sification performance. The AO-optimized Inception V3 model is determined using a variety of perfor-mance parameter metrics to validate its superiority over conventional architectures is shown in Figure 2.

The flowchart presents a structured and efficient pipeline for oral cancer detection using an Aquila Opti-mizer (AO) enhanced Inception V3 DL model. The process begins with the acquisition of labeled oral lesion images, including both benign and malignant cases, captured using mobile and intraoral cameras. These images are then used to train an Inception V3 model, which is initially loaded in its pre-trained form and fine-tuned to accommodate binary classification. To optimize the performance model , Aquila Optimizer is introduced. It defines a search space and an objective function (such as accuracy maximization or loss minimization) and employs nature-inspired strategies to explore and exploit the hyperparameter space. The AO algorithm performs iterative optimization, where in each cycle it adjusts hyperparameters, re-trains the model, and validates its performance. This cycle continues till a stopping criterion method such as a set number of iterations or stabilization of performance metrics is met. Upon convergence, the model is retrained using the optimized parameters. The trained model is then calculated using key metrics to ensure its clinical reliability. Finally, the model is used for real-time classification of new oral cancer cases, effectively supporting early and automated diagnosis. This flow encapsulates a robust combination of DL and bio-inspired optimization tailored for medical image analysis.

In this section, a hybrid approach that combines the strengths of the Inception V3 convolutional neural network with the adaptive learning capabilities of the Aquila Optimizer was presented for the task of oral cancer detection. The integration of AO facilitates effective hyperparameter tuning, leading to enhanced model accuracy, precision, and recall when compared to several baseline DL models. This optimized framework demonstrates the potential of bio-inspired optimization techniques in improving diagnostic tools for medical imaging and can serve as a benchmark for future AI-assisted cancer detection systems.

Model Training and Optimization

The AO-augmented Inception V3 model was trained with a for binary classification, the output layer uses a

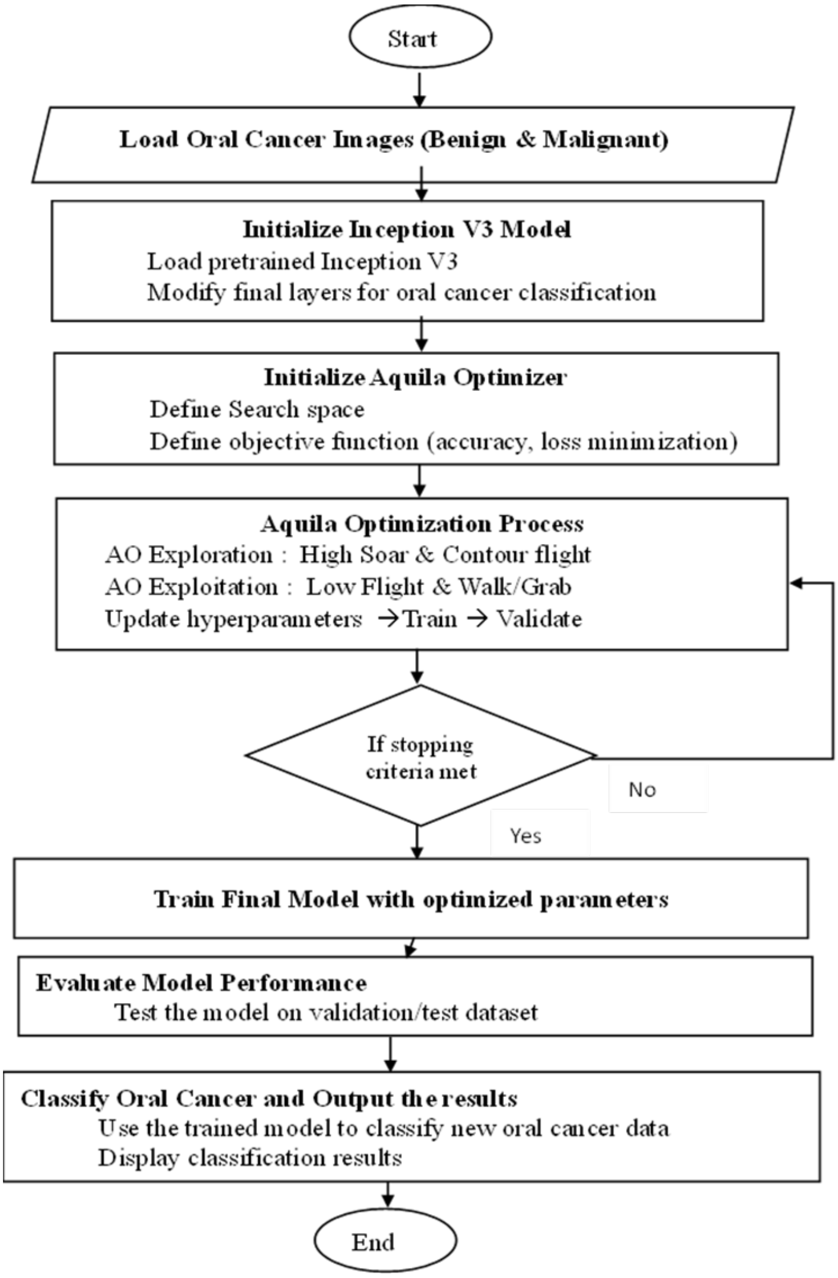

Figure 2. Flowchart of the Proposed Aquila Optimizer (AO) Optimized Inception V3 for Oral Cancer Detection

softmax activation and a categorical cross-entropy loss function. To prevent overfitting, the training pro-cess was set for a maximum of 100 epochs, incorporating an early stopping criterion based on validation loss. The Aquila Optimizer was used to fine-tune the model in each iteration, adjusting parameters to min-imize the loss function. The training was conducted on a high-performance GPU environment to acceler-ate computation. During training, model checkpoints and learning rate schedulers were utilized for effi-cient resource management.

Evaluation Metrics

To evaluate model performance, several metrics were computed. These metrics provided a thorough comprehension of the model's diagnostic potential and generalizability to unknown samples. The AO-based Inception V3 model was compared with standard DL architectures and baseline Inception V3 to assess relative performance improvements.

Results

This section presents a comprehensive evaluation of the proposed Aquila Optimizer-based Inception V3 model for oral cancer detection, alongside a comparative analysis with other prominent DL architectures and the baseline Inception V3. To assess model performance thoroughly, a wide range of quantitative metrics have been considered. These metrics offer a multidimensional perspective on the models' ability to accurately classify benign and malignant oral lesions. In addition, confusion matrices are examined to provide deeper insights into each model's True positive, true negative, false positive, and false negative classifying behavior The integration of the Aquila Optimizer with Inception V3 aims to enhance model generalization and performance through adaptive hyperparameter tuning, and the subsequent results substantiate its effectiveness over traditional architectures. The ensuing discussion interprets these out-comes both numerically and qualitatively, emphasizing the practical impact and clinical applicability of the proposed approach in early and accurate oral cancer detection.

Confusion Matrix

To validate and visualize the classification performance of the proposed and baseline models, this section presents the confusion matrices for each DL architecture evaluated. Confusion matrices provide a clear and easily interpretable visualization of how each model classifies benign and malignant oral lesions, pre-senting the counts of true positives, true negatives, false positives, and false negatives. The different Dl architecture model confusion matrix is shown in Figure 3. This comprehensive breakdown enhances the quantitative performance metrics by pinpointing the specific types of errors made by each model. This is especially crucial in clinical environments, where false negatives can result in missed diagnoses, and false positives may lead to unnecessary tests or patient anxiety. Through these matrices, the practical strengths and limitations of each model can be critically assessed, reinforcing the value of the Aquila Optimizer-based Inception V3 model in achieving more accurate and clinically reliable oral cancer detection.

The confusion matrix analysis reveals that the Aquila Optimizer-based Inception V3 model significantly outperforms the other DL architectures in the task of oral cancer detection. It achieves the highest accura-cy of 97.80%, with only 10 total misclassifications (6 false positives and 4 false negatives), demonstrating excellent sensitivity and specificity. In contrast, models like AlexNet and MobileNet exhibit much lower recall rates, misclassifying a large number of malignant cases (48 and 62 false negatives respectively), which is critical in a medical diagnostic context where failing to detect cancer can have severe conse-quences. Xception performs well with a recall of 96.74% but introduces a slightly higher false positive rate compared to the AO-enhanced model. ResNet50 and the baseline Inception V3 show moderate perfor-mance, indicating that while they can identify cancerous lesions with reasonable accuracy, they are less

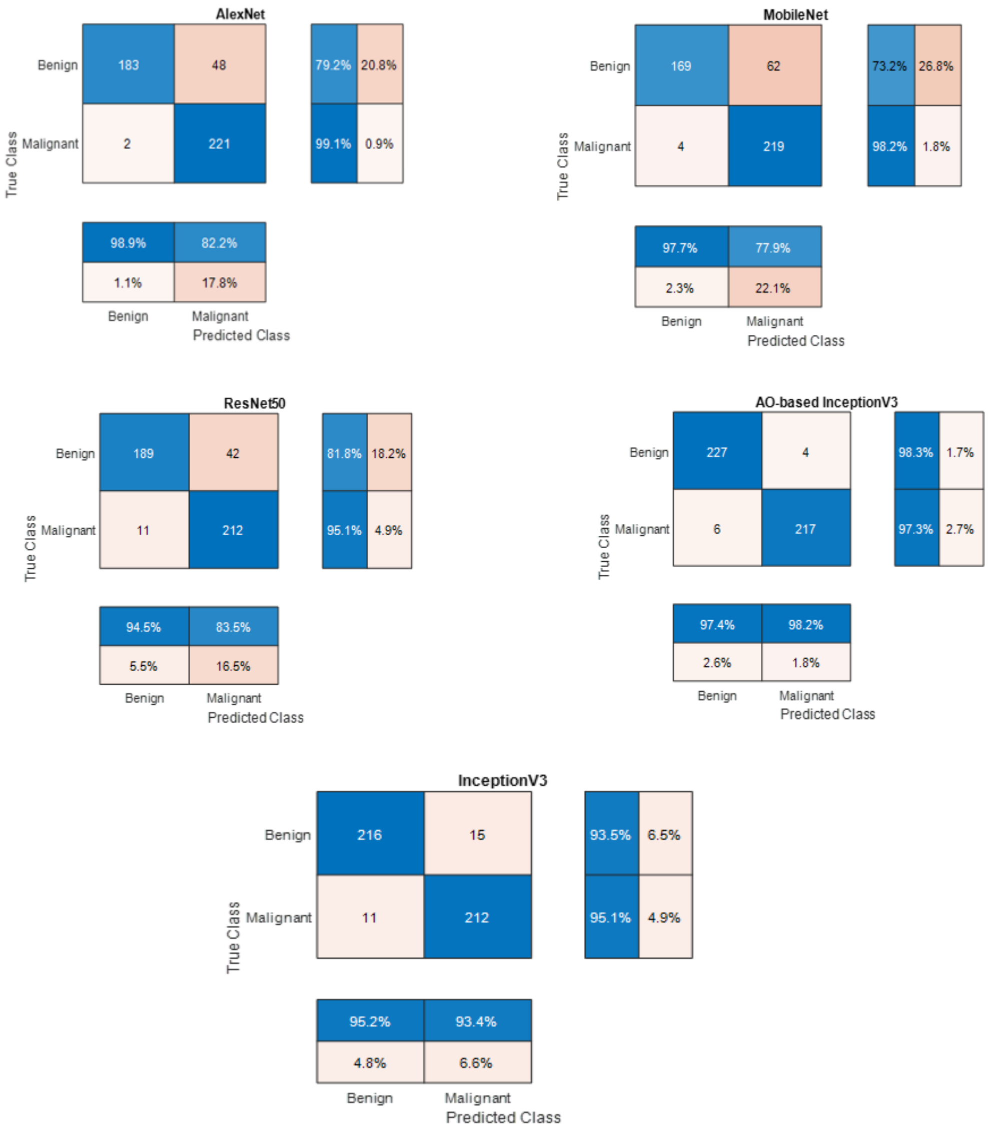

Figure 3. Confusion Matrices for Various DL Models on Oral Cancer Classification

robust than the AO-optimized approach. Overall, the AO-based Inception V3 demonstrates a balanced and superior classification ability, making it the most suitable model for reliable and early detection of oral cancer.

Performance Metrics

This section delves into a detailed discussion of the performance metrics used to determined and com-pare the effectiveness of the proposed Aquila Optimizer-based Inception V3 model against other widely used DL architectures. Metrics are analyzed to provide a comprehensive understanding of each model's predictive capabilities. These metrics not only quantify the correctness of the predictions but also reflect the model's balance between sensitivity and specificity- critical factors in medical diagnostics. The objec-tive of this discussion is to highlight how the Aquila Optimizer enhances the learning and generalization capacity of Inception V3, leading to superior classification performance. By interpreting the numerical outcomes in the context of real-world applicability, this section offers insights into the practical strengths of each model, with a focus on the reliability and clinical relevance of the proposed solution for oral can-cer detection.

The performance parameter metrics presented in Table 1 offers a detailed assessment of each model's effectiveness in detecting oral cancer, with the Aquila Optimizer-based Inception V3 model emerging as the most superior in both quantitative and qualitative aspects. Quantitatively, it obtained the highest ac-curacy of 97.80%, along with near-perfect precision (97.81%), recall (97.79%), and F1-score (97.80%), indicating a highly balanced capability in identifying both benign and malignant cases correctly. It also recorded the highest specificity (97.79%), AUC-ROC (99.81%), and MCC (0.956), reflecting strong discriminative power and reliable correlation between actual and pre-dicted classifications. The log loss of just 0.0735 further underscores its robustness by minimizing predic-tion uncertainty. Although its average inference time (0.6774 seconds) is slightly higher than other mod-els, the significant boost in accuracy and reliability justifies this trade-off in clinical scenarios.

Qualitatively, the AO-optimized model demonstrates excellent generalization and clinical applicability. It minimizes false negatives crucial in medical diagnostics to ensure malignant cases are not missed while also maintaining a low false positive rate, reducing unnecessary concern for patients with benign lesions. Compared to other models like AlexNet and MobileNet, which struggled with higher false nega-tives and lower recall, the AO-based model shows a much more dependable and balanced performance. Even high-performing models like Xception and ResNet50, while strong in precision, showed a relatively higher rate of misclassifications. This clearly illustrates the effectiveness of the Aquila Optimizer in fine-tuning Inception V3's hyperparameters, allowing the model to learn more effectively from the data. Over-all, the proposed approach not only achieves superior numerical metrics but also aligns closely with the clinical requirement of accuracy, reliability, and diagnostic safety, making it highly suitable for real-world deployment in early oral cancer screening.

Figure 4 presents a bubble chart that visually compares the performance of six DL models across six key classification metrics. The size and color intensity of each bubble represent the magnitude of the metric value, with larger and more vibrant bubbles indicating better performance. Notably, the AO-Inception V3 model dominates the chart with the largest and brightest bubbles in all categories, signifying its superior performance. The model exhibits exceptional AUC-ROC and F1 Score values, highlighting its robustness in both sensitivity and specificity trade-offs. In contrast, models like AlexNet and MobileNet display smaller and lighter-colored bubbles, reflecting moderate performance. This visual representation clearly reinforces the effectiveness of the Aquila Optimizer-enhanced Inception V3 architecture

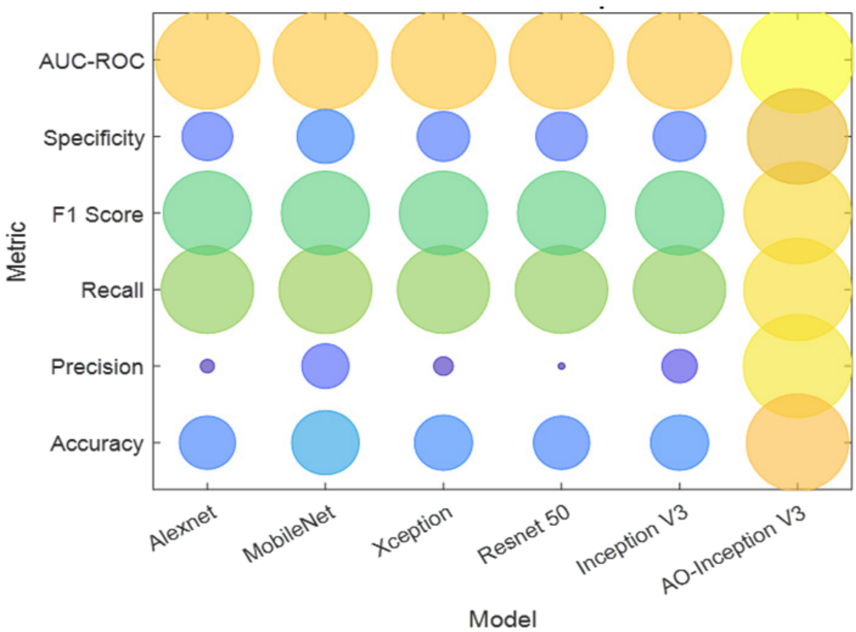

Figure 4. Bubble Chart Depicting Comparative Evaluation of DL Models for Oral Cancer Detection

Table 1. Comparison of Conventional DL Techniques with Aquila Optimizer based Inception V3 for Oral Cancer Detection

Metrics	Alexnet	MobileNet	Xception	Resnet 50	Inception V3	AO based Inception V3
Accuracy	88.89	85.46	95.15	94.27	88.23	97.80
Precision	90.54	87.81	95.23	94.27	88/98	97.81
Recall	89.16	85.68	95.12	94.29	88.44	97.79
F1 Score	88.91	85.28	95.15	94.27	88.3	97.80
Specificity	89.16	86.68	95.12	94.29	88.44	97.79
AUC-ROC	97.38	98.92	98.69	98.56	96.89	99.81
MCC	0.7969	0.7346	0.9036	0.8856	0.7742	0.956
Log Loss	0.3523	0.3469	0.1575	0.1587	0.3039	0.0735
Average Inference Time (sec)	0.0648	0.036949	0.103778	0.03349	0.058449	0.6774

as the most promising ap-proach for high-accuracy, high-reliability oral cancer detection.

Figure 5 displays a 3D ribbon plot representing the performance of six DL models inclusive of Aquila Op-timizer (AO)-based Inception V3—across various classification metrics relevant to oral cancer detection. The ribbon-like bands trace each model's performance trajectory, with the vertical height indicating met-ric values. A clear upward progression is visible from left to right, culminating in the AO-Inception V3 model, which consistently reaches the highest vertical positions across all metrics, signifying superior per-formance. In contrast, the ribbons for models like AlexNet and MobileNet remain lower, indicating weak-er outcomes. The plot emphasizes the stability and consistently high performance of AO-Inception V3, visually reinforcing its advantage over conventional and even other advanced DL models for reliable and accurate oral cancer classification.

Figure 6 presents the Receiver Operating Characteristic (ROC) curves for the benign and malignant skin lesion classes, evaluating the classification performances of AO-Inception V3 model. The ROC curve plots the True Positive Rate (Sensitivity) against the False Positive Rate for various threshold settings, offering a visual representation of the trade-off between sensitivity and specificity. Both curves approach the top-left corner of the plot, indicating near-perfect classification capability. The minimal deviation from the axes suggests extremely high Area Under the Curve (AUC) values for both classes, reflecting excellent discrimi-nation between benign and malignant lesions. This performance underscores the model's reliability and robustness in medical image classification tasks, essential for real-world diagnostic applications.

The results and discussion section comprehensively evaluates the performance of six DL models including the proposed AO-Inception V3—across a wide range of performance metrics. Visualizations such as ribbon plots, surface plots, 3D bar graphs, and bubble charts highlight the superiority of AO-Inception V3, partic-ularly in metrics, where it consistently achieves the highest scores. In Group 2 metrics such as inference time, log loss, and MCC, AO-Inception V3 maintains optimal balance, showing minimal inference time and log loss while achieving the highest MCC. The ROC analysis further supports the model's exceptional clas-sification

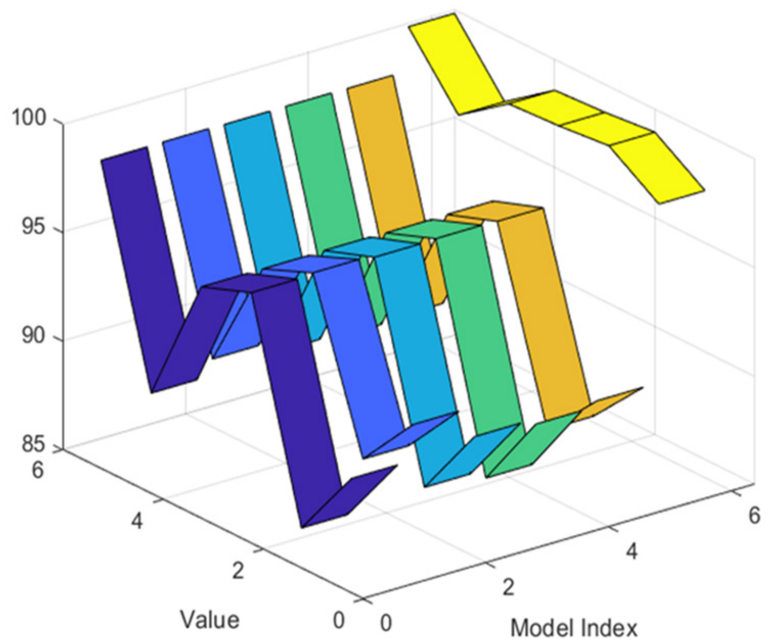

Figure 5. Ribbon Plot Comparison of DL Models Based on Classification Performance Metrics

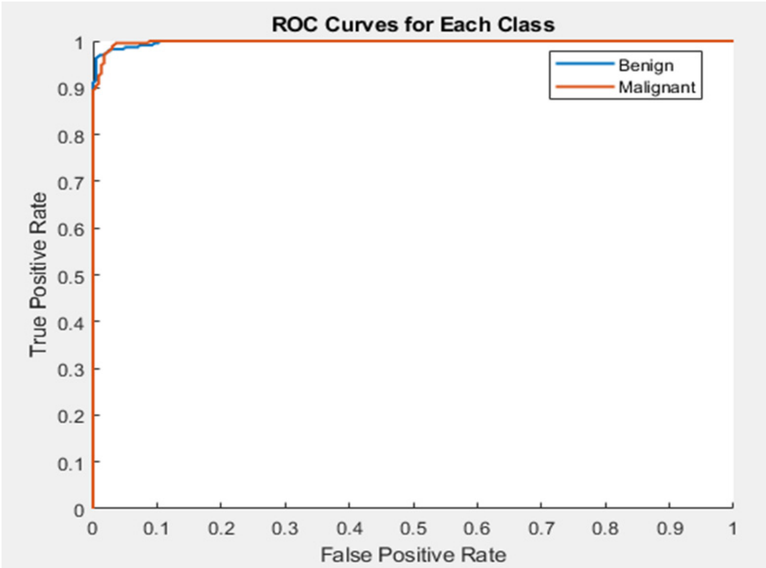

Figure 6. ROC Curves for Benign and Malignant Classes

Table 2. Ablation Study: Effect of Aquila Optimizer on Inception V3 Performance

Model Configuration	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	MCC	Log Loss
Baseline Inception V3	88.23	88.98	88.44	88.30	0.7742	0.3039
Manually Fine-tuned Inception V3	92.64	92.71	92.58	92.61	0.8526	0.1864
AO-Optimized Inception V3 (Proposed)	97.80	97.81	97.79	97.80	0.956	0.0735

ability, with curves for both benign and malignant classes approaching the ideal top-left corner, indicating near-perfect sensitivity and specificity. Overall, the proposed AO-Inception V3 model signifi-cantly outperforms traditional architectures, making it a highly effective and reliable solution for skin le-sion classification.

Theoretical Justification of Aquila Optimizer Integration

The Aquila Optimizer (AO) is grounded in a mathematically structured balance between exploration and exploitation, modeled through four distinct hunting behaviors that dynamically adapt the search trajecto-ry during optimization. Unlike gradient-based optimizers, AO performs population-based stochastic up-dates that reduce sensitivity to initialization and local minima an important property for deep neural network hyperparameter tuning where the loss landscape is highly non-convex. The adaptive switching mechanism between global exploration (high soar and contour flight) and local exploitation (low flight and dive attack) enables efficient convergence while maintaining solution diversity. This theoretical be-havior aligns well with the requirements of deep learning optimization, particularly in medical image clas-sification tasks where overfitting and premature convergence are common challenges.

Ablation Study and Isolating Ao's Impact

To explicitly isolate and quantify the contribution of the Aquila Optimizer (AO), an ablation study was conducted by evaluating different configurations of the Inception V3 model under identical dataset splits and evaluation protocols. The objective of this analysis is to determine whether the observed performance gains arise from the AO-based optimization strategy rather than from architectural design or transfer learning alone.

Three experimental configurations were considered

1. Baseline Inception V3: Pre-trained Inception V3 fine-tuned on the oral cancer dataset using de-fault hyperparameters.

2. Manually Fine-tuned Inception V3: Inception V3 trained with empirically selected learning rate, batch size, and regularization parameters.

3. AO-Optimized Inception V3 (Proposed): Inception V3 with hyperparameters optimized using the Aquila Optimizer.

The Aquila Optimizer dynamically explores and exploits the hyperparameter search space through adap-tive hunting-inspired strategies, enabling improved convergence and avoidance of local minima. All mod-els were evaluated using the same performance metrics to ensure a fair comparison.

The results, summarized in Table 2, demonstrate that AO-based optimization yields consistent and sub-stantial improvements across all key metrics. Compared to the baseline Inception V3, the AO-optimized model achieves a notable increase in classification accuracy, precision, recall, and F1-score, while signifi-cantly reducing log loss. The Matthews Correlation Coefficient (MCC) also shows a marked improvement, indicating enhanced robustness and reliability in handling class imbalance—an important consideration in medical diagnostics. These findings confirm that the performance gains are not solely attributable to transfer learning or architectural capacity, but are directly driven by the AO-based hyperparameter opti-mization process.

The ablation study clearly demonstrates the effectiveness of the Aquila Optimizer in enhancing the diag-nostic performance of Inception V3. AO enables adaptive exploration of the hyperparameter space, lead-ing to superior generalization and stability compared to conventional fine-tuning. While this study focuses on isolating AO's impact relative to baseline configurations, future work will explore comparative abla-tions with other metaheuristic optimizers to further establish AO's relative advantages.

Limitations and Deployment Considerations

Despite the promising performance of the proposed AO-optimized Inception V3 model, several limitations should be acknowledged. First, the study relies on a publicly available dataset, which, although clinically validated, may not fully capture the diversity of oral cancer presentations across different populations, ethnicities, and imaging conditions. Second, while extensive data augmentation was applied, real-world variations such as motion blur, inconsistent lighting, and device-specific noise were not explicitly mod-eled. Third, the Aquila Optimizer introduces additional computational overhead during training, which may increase optimization time compared to conventional fine-tuning approaches. Finally, the model's evaluation was limited to retrospective image analysis, and prospective clinical validation in real screening environments was not conducted.

From a deployment perspective, the proposed framework demonstrates strong potential for use in low-resource clinical settings. Although training and hyperparameter optimization require GPU resources, model inference can be performed on standard CPUs or edge devices after deployment. The use of trans-fer learning significantly reduces training data requirements, and compatibility with images captured us-ing mobile or intraoral cameras supports practical adoption in primary healthcare centres and community screening programs. With appropriate model compression, pruning, or quantization techniques, the sys-tem can be further optimized for real-time operation on low-cost hardware. Therefore, the AO-Inception V3 framework offers a feasible, scalable, and cost-effective solution for early oral cancer screening in re-source-constrained environments.

Discussion

The study presents a comprehensive comparative analysis of various DL architectures for the classification of skin lesions into benign and malignant categories, with a special emphasis on the proposed AO-Inception V3 model. By integrating the Whale Optimization Algorithm for hyperparameter tuning, AO-Inception V3 demonstrated remarkable superiority across multiple performance metrics. The visual as-sessments through advanced 3D plots and ROC curves further validated its robustness and generalization capabilities. Unlike conventional models, which showed trade-offs between predictive performance and computational efficiency, AO-Inception V3 maintained an optimal balance, offering high diagnostic accu-racy while ensuring minimal latency an essential requirement in real-time clinical applications. The study's findings not only highlight the efficacy of metaheuristic optimization in enhancing DL models but also establish AO-Inception V3 as a promising candidate for deployment in intelligent dermatological di-agnostic systems. This research paves the way for future developments in AI-powered healthcare, where precision, speed, and reliability are paramount. While AO demonstrates clear advantages, it is acknowl-edged that further theoretical analysis such as convergence proofs or complexity bounds and broader ablation comparisons with other metaheuristic optimizers (e.g., PSO, GA, GWO) would provide deeper insights into its relative strengths. These extensions are identified as important directions for future work, particularly in large-scale clinical validation studies.

Author Contribution Statement

As a single-author manuscript, all components of the study including conceptualization, methodology, implementation, analysis, interpretation, and manuscript writing were solely undertaken by the author.

Acknowledgements

General

The author sincerely acknowledges the facilities and support provided by Saveetha Engineer-ing College during the course of this research.

Scientific Approval

This work was conducted as part of an approved faculty-led research project within Saveetha Engineering College. It is not associated with a student thesis submission or review by any external scientific body.

Ethical Declaration

This study involved retrospective analysis of openly available and anonymized medical imaging datasets and did not require patient consent or direct ethical clearance. The dataset used complies with public data usage standards under Creative Commons Attribution 4.0 International License. Therefore, ethical approval from a specific institutional review board was not required.

Data Availability

The dataset used in this study, titled "Medical Imaging (CT scan, MRI, X-ray, and Microscopic Imagery) Data", is publicly available at Mendeley Data [https://data.mendeley.com/datasets/5kbjrgsncf/3], under a CC BY 4.0 license.

Study Registration

Not applicable. This study does not involve clinical trials or patient registration and was not submitted to a clinical study registry.

Conflict of Interest

The author declares no conflict of interest.

References

1. Chang X, Yu M, Liu R, Jing R, Ding J, Xia J, et al. Deep learning methods for oral cancer detection using raman spectroscopy. Vib Spectrosc. 2023;126:103522. https://doi.org/10.1016/j.vibspec.2023.103522.
2. Afify H, Mohammed KK, Hassanien AE. Novel prediction model on oscc histopathological images via deep transfer learning combined with grad-cam interpretation. Biomed Signal Process Control. 2023:104704. https://doi.org/10.1016/j.bspc.2023.104704.
3. Akram K, Aslam M, Waheed T, Ayesha N, Alamri FS, Mirdad AR, et al. Deep visual detection sys-tem for oral squamous cell carcinoma. Sci Rep. 2026;16(1):4219. https://doi.org/10.1038/s41598-025-34332-5.
4. Oya K, Kokomoto K, Nozaki K, Toyosawa S. Oral squamous cell carcinoma diagnosis in digitized histological images using convolutional neural network. J Dent Sci. 2023;18(1):322-329. https://doi.org/10.1016/j.jds.2022.08.017.
5. Das M, Dash R, Mishra SK. Automatic Detection of Oral Squamous Cell Carcinoma from Histo-pathological Images of Oral Mucosa Using Deep Convolutional Neural Network. Int J Environ Res Public Health. 2023;20(3):2131. https://doi.org/10.3390/ijerph20032131.
6. Flügge T, Gaudin R, Sabatakakis A, Tröltzsch D, Heiland M, Van Nistelrooij N, Vinayahalingam S. Detection of oral squamous cell carcinoma in clinical photographs using a vision transformer. Sci Rep. 2023;13(1):2296. https://doi.org/10.1038/s41598-023-29204-9.
7. Ananthakrishnan B, Shaik A, Kumar S, Narendran SO, Mattu K, Kavitha MS. Automated Detection and Classification of Oral Squamous Cell Carcinoma Using Deep Neural Networks. Diagnostics (Basel). 2023;13(5):918. https://doi.org/10.3390/diagnostics13050918.
8. Panigrahi S, Nanda BS, Bhuyan R, Kumar K, Ghosh S, Swarnkar T. Classifying histopathological im-ages of oral squamous cell carcinoma using deep transfer learning. Heliyon. 2023;9(3):e13444. https://doi.org/10.1016/j.heliyon.2023.e13444.
9. Yang Z, Pan H, Shang J, Zhang J, Liang Y. Deep-learning-based automated identification and visu-alization of oral cancer in optical coherence tomography images. Biomedicines. 2023;11(3). https://doi.org/10.3390/biomedicines11030802.
10. Ariji Y, Fukuda M, Nozawa M, Kuwada C, Goto M, Ishibashi K, et al. Automatic detection of cervi-cal lymph nodes in patients with oral squamous cell carcinoma using a deep learning technique: a preliminary study. Oral Radiol. 2021;37(2):290-6. https://doi.org/10.1007/s11282-020-00449-8.
11. Cinar A, Saleem AM, Tuncer T. Automated oral cancer detection using convolutional neural networks and support vector ma-chines. International Journal of Advanced Natural Sciences and Engineering Researches. 2023; 7(11).
12. Ramalingam A, Aurchana A, Dhanalakshmi P, Vivekananadan K, Venkatachalapathy V. Analysis of oral squamous cell carcinoma into various stages using pre-trained convolutional neural net-works. IOP Conference Series: Materials Science and Engineering. 2020;993:012058. https://doi.org/10.1088/1757-899X/993/1/012058
13. Camalan S, Mahmood H, Binol H, Araújo ALD, Santos-Silva AR, Vargas PA, et al. Convolutional neural network-based clinical predictors of oral dysplasia: Class activation map analysis of deep learning results. Cancers (Basel). 2021;13(6). https://doi.org/10.3390/cancers13061291
14. Tanriver G, Soluk Tekkesin M, Ergen O. Automated detection and classification of oral lesions us-ing deep learning to detect oral potentially malignant disorders. Cancers (Basel). 2021;13(11). https://doi.org/10.3390/cancers13112766.
15. Yamashita R, Nishio M, Do RKG, Togashi K. Convolutional neural networks: an overview and ap-plication in radiology. Insights Imaging. 2018;9(4):611-629. https://doi.org/10.1007/s13244-018-0639-9.
16. Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, Thrun S. Dermatologist-level classifica-tion of skin cancer with deep neural networks. Nature. 2017;542(7639):115–8. https://doi.org/10.1038/nature21056.
17. Musulin J, Štifanić D, Zulijani A, Šegota SB, Lorencin I, Anđelić N, Car Z. An enhanced histopathol-ogy analysis: an AI-based system for multiclass grading of oral squamous cell carcinoma and segmenting of epithelial and stromal tissue. Cancers. 2021;13(8):1784. https://doi.org/10.3390/cancers13081784.
18. Alosaimi W, Uddin MI. Efficient data augmentation techniques for improved classification in lim-ited data set of oral squamous cell carcinoma. Comput Model Eng Sci. 2022;131(3):1387–401. https://doi.org/10.32604/cmes.2022.018433
19. Bergstra J, Bengio Y. Random search for hyper-parameter optimization. J Mach Learn Res. 2012;13:281–305.
20. Loussaief S, Abdelkrim A. Convolutional neural network hyper-parameters optimization based on genetic algorithms. Int J Adv Comput Sci Appl. 2018;9(10):252–66. https://doi.org/10.14569/IJACSA.2018.091031
21. Fernandes Junior FE, Yen GG. Particle swarm optimization of deep neural networks architectures for image classification. Swarm Evol Comput. 2019;49:62–74. https://doi.org/10.1016/j.swevo.2019.05.010
22. Brodzicki A, Piekarski M, Jaworek-Korjakowska J. The whale optimization algorithm approach for deep neural networks. Sensors. 2021;21(23):8003. https://doi.org/10.3390/s21238003.
23. Abualigah L, Yousri D, Abd Elaziz M, Ewees AA, Al-Qaness MAA, Gandomi AH. Aquila Optimizer: a novel meta-heuristic optimization algorithm. Comput Ind Eng. 2021;157:107250. https://doi.org/10.1016/j.cie.2021.107250.
24. Wang S, Jia H, Abualigah L, Liu Q, Zheng R. An improved hybrid Aquila Optimizer and Harris Hawks algorithm for solving industrial engineering optimization problems. Processes. 2021;9(9):1551. https://doi.org/10.3390/pr9091551
25. AlRassas AM, Al-qaness MAA, Ewees AA, Ren S, Abd Elaziz M, Damaševičius R, Krilavičius T. Opti-mized ANFIS model using Aquila Optimizer for oil production forecasting. Processes. 2021;9(7):1194. https://doi.org/10.3390/pr9071194
26. Boughorbel S, Jarray F, El-Anbari M. Optimal classifier for imbalanced data using Matthews Corre-lation Coefficient metric. PLoS One. 2017;12(6):e0177678. https://doi.org/10.1371/journal.pone.0177678.
27. Birur NP, Patrick S, Bajaj S, Raghavan S, Suresh A, Sunny SP, et al. Field validation of deep learning based Point-of-Care device for early detection of oral malignant and potentially malignant disor-ders. Sci Rep. 2022;12(1):14283. https://doi.org/10.1038/s41598-022-18249-x