Hybrid Metaheuristic-Optimized Deep Learning for Interpretable and Fair Early Detection of Oral Squamous Cell Carcinoma: A Systematic Review and Methodological Framework

Khaled Sh. Gaber[1,*], **Shahid Mahmood**[2]
[1]Computer Science and Intelligent Systems Research Center, Blacksburg 24060, Virginia, USA
[2]School of Finance and Economics, Jiangsu University, Zhenjiang, People's Republic of China
Emails: khsherif@jcsis.org; shahidnajam786@live.com

Abstract

Oral cancer remains a significant global health concern, particularly due to the high rates of late-stage detection and the limitations of traditional diagnostic modalities. This study proposes a hybrid diagnostic framework that integrates deep learning with metaheuristic optimization to enhance the accuracy, efficiency, and interpretability of oral cancer classification. The architecture combines convolutional and recurrent neural network components with an adaptive optimization layer designed using swarm intelligence-inspired algorithms. This hybridization enables precise feature selection, architecture tuning, and parameter optimization, resulting in improved generalization and robustness across heterogeneous clinical datasets. The model is further augmented with explainable decision support features, enabling clinicians to visualize lesion relevance and interpret classification outcomes. Empirical evaluations demonstrate superior performance in terms of sensitivity, specificity, and computational efficiency compared to conventional training strategies. Additionally, the proposed framework is designed for portability and scalability, supporting potential deployment in mobile and edge-based diagnostic systems. The integration of interpretability, fairness constraints, and clinical adaptability underscores the model's readiness for real-world implementation. This work contributes to the growing field of intelligent medical diagnostics and highlights the transformative potential of metaheuristic optimization in addressing complex, high-dimensional clinical classification tasks.

Keywords: Metaheuristic optimization; Deep learning; Feature selection; Oral cancer; Intelligent diagnostics

1 Introduction

Oral cancer, particularly oral squamous cell carcinoma (OSCC), constitutes a pressing global health burden and continues to be a focus of intense clinical and research interest due to its high prevalence, aggressive progression, and complex pathology. It remains one of the most commonly diagnosed malignancies in the head and neck region and is frequently associated with a high degree of morbidity and mortality, especially when detection is delayed or treatment is initiated in the advanced stages of the disease. Despite significant advancements in imaging technologies, molecular profiling, and treatment strategies including surgery, chemotherapy, and radiotherapy, clinical outcomes for patients diagnosed at later stages remain suboptimal. The five-year survival rate remains unsatisfactory in many regions, particularly in low- and middle-income countries, where access to advanced diagnostic tools and early detection programs is limited. A key limitation in the contemporary diagnostic pipeline for oral cancer lies in the reliance on invasive, time-consuming, and often subjective

DOI: https://doi.org/10.54216/MOR.050203
Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

procedures such as surgical biopsies, histopathological interpretation, and manual image analysis. These approaches, while considered the gold standard, suffer from inter-observer variability, delays in diagnosis, and limited scalability in population-wide screening efforts. Consequently, there is a growing consensus in the medical informatics and computational biology communities on the need for innovative, non-invasive, and scalable computational solutions capable of augmenting or even partially automating the diagnostic workflow. In response to these challenges, the emergence and rapid evolution of artificial intelligence (AI), machine learning (ML), and particularly metaheuristic optimization techniques have opened new pathways in the development of intelligent diagnostic systems. Metaheuristic algorithms—an umbrella term encompassing a variety of nature-inspired, stochastic, and population-based search methods—have gained substantial traction in biomedical research due to their ability to efficiently solve complex, high-dimensional, and non-convex optimization problems that are otherwise intractable using classical deterministic methods. These algorithms mimic processes found in nature, such as evolutionary selection, swarm cooperation, or physical dynamics, and are capable of balancing global exploration and local exploitation to identify optimal or near-optimal solutions within large and noisy solution spaces. A particularly well-explored application of metaheuristics in the medical domain is in the optimization of support vector machine (SVM) classifiers, which are widely used in pattern recognition tasks due to their strong generalization ability and solid theoretical foundation. However, the performance of SVMs is highly dependent on the appropriate selection of hyperparameters such as kernel type, cost penalty, and gamma values. The use of metaheuristics for automatic parameter tuning has demonstrated considerable success in enhancing the classification accuracy of SVMs in biomedical tasks. A notable example is demonstrated in the optimization of SVMs for the differential diagnosis and classification of oral diseases using clinical image datasets containing multiple lesion types, including leukoplakia, lichen planus, and OSCC [1]. Extending beyond classical ML techniques, recent research has focused on the integration of metaheuristics with deep learning frameworks. Deep neural networks, while powerful in modeling complex, hierarchical data, suffer from issues related to overfitting, local minima, and computational cost during training. Metaheuristic optimization methods have been employed to address these limitations by automating the selection of network architecture parameters, activation functions, learning rates, and dropout ratios. This has resulted in more efficient, accurate, and stable deep learning models tailored for oral cancer detection. One approach involves metaheuristic-based deep feature selection mechanisms, which enable the network to focus only on the most discriminative and clinically relevant features extracted from oral lesion images, thus reducing dimensionality while preserving diagnostic integrity [2]. Further enhancements to model performance and convergence have been achieved through hybrid learning architectures. One such architecture involves the fusion of recurrent and belief-based networks trained using a custom metaheuristic optimization strategy. In such frameworks, learning parameters—including the number of hidden neurons, learning rate, and the number of training epochs—are tuned using advanced swarm intelligence algorithms that combine the behavior of biological entities such as beetles and barnacles to achieve high-performance classification with rapid convergence and noise resilience. These hybrid models have been shown to surpass traditional neural architectures in terms of accuracy, sensitivity, and overall training efficiency when applied to noisy or low-resolution image datasets often encountered in oral cancer research [3]. To further overcome the limitations of singular optimization methods, recent methodologies have introduced composite metaheuristic frameworks that integrate the strengths of multiple algorithms. An example is the combination of particle swarm optimization with earth-radius-based geospatial heuristics, resulting in an algorithm capable of achieving improved convergence, stability, and classification precision, especially in challenging oral cancer datasets characterized by high intra-class variability and low inter-class separability. These hybrid optimization strategies are particularly beneficial in fine-tuning deep learning classifiers and ensuring robustness against local optima. Early-stage detection, a critical factor in improving oral cancer prognosis, has also been a primary focus in AI-driven diagnostic system development. To this end, innovative combinations of deep belief networks and group-based teaching optimization strategies have been utilized to improve classification performance at the early detection stage. Such frameworks emphasize collaborative learning and knowledge-sharing among model agents to accelerate convergence and reduce the training time while maintaining high classification fidelity in distinguishing early-stage malignancies from benign and pre-malignant lesions [4]. The use of convolutional neural networks (CNNs) remains a cornerstone of image-based cancer diagnostics. Optimization techniques have been successfully applied to CNN training, specifically in the context of oral lesion detection, where convolutional layer configurations, kernel sizes, and pooling strategies are adaptively tuned using seagull-inspired metaheuristics. These biologically inspired optimization processes ensure that CNNs are not only effective in feature learning but also efficient in training, contributing to their suitability for deployment in clinical and mobile health settings [5]. Similarly, the integration of recurrent units optimized through bird-of-prey-inspired algorithms has yielded architectures with enhanced temporal modeling capabilities. These models are particularly valuable in clinical scenarios involving sequential image capture or lesion progression

DOI: https://doi.org/10.54216/MOR.050203

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

tracking over time. The optimization of such recurrent structures ensures stability and adaptability in recognizing complex lesion growth patterns and underlying malignancy trends [6]. In parallel to these deep learning strategies, metaheuristic-driven frameworks have also been applied successfully to other cancer types and modalities, demonstrating cross-domain applicability and transferability of the optimization paradigms. For instance, diagnostic models developed for cervical cytology images using metaheuristic optimization strategies and deep learning pipelines reveal the versatility of these techniques and inform their adaptation to oral cancer pathology [7]. Advances in real-time diagnostic systems have also been enabled by the integration of metaheuristically optimized models into IoT-based healthcare frameworks. These systems are designed to operate within constrained computational environments and provide immediate diagnostic feedback, thereby enhancing accessibility and speed in oral cancer screening. The use of lightweight deep learning models optimized by efficient heuristics allows such systems to perform effectively even on embedded or mobile hardware [8]. Further innovations involve the use of aquatic-life-inspired optimization techniques integrated with CNN architectures. These optimization strategies enhance the ability of CNNs to extract complex features from medical images and improve classification outcomes, particularly in detecting early oral cancers and differentiating them from benign lesions [9]. The application of sine cosine algorithm variants to fine-tune deep learning model parameters has similarly contributed to increased training efficiency and model accuracy in oral cancer diagnosis tasks [10]. In recognition of the need for comprehensive reviews and critical evaluations of these evolving methodologies, recent literature has synthesized findings from a broad range of studies to highlight emerging trends, challenges, and future directions in the application of deep learning and optimization in oral cancer diagnostics. These reviews underscore the growing emphasis on interpretability, model efficiency, and integration with clinical workflows [11]. Finally, advanced architectures that leverage attention mechanisms and feature fusion strategies have demonstrated remarkable improvements in oral cancer detection. These models dynamically prioritize the most relevant image regions and integrate multi-scale contextual information to enhance classification performance, enabled in part by sophisticated metaheuristic strategies for parameter tuning and architectural optimization [12]. Taken together, these diverse yet thematically convergent developments illustrate the transformative role of metaheuristic optimization in advancing oral cancer diagnosis. By enabling Building upon the foundational discussion outlined earlier, further elaboration is necessary to contextualize the broader significance of the present study's findings within the dynamic and multidisciplinary field of computational oncology. While the hybrid metaheuristic-driven framework developed herein has shown substantial promise in enhancing classification accuracy, interpretability, and model robustness, a deeper analytical dissection reveals several pivotal layers of methodological, clinical, and translational impact. One of the key contributions of this research lies in the application of biologically inspired optimization strategies to deep learning-based oral cancer diagnostic systems. Unlike conventional gradient-based optimizers, which often exhibit brittle convergence behaviors in the presence of non-stationary data distributions or sparse minority-class representation, the use of swarm and evolutionary heuristics enables more flexible exploration of the model configuration space. This flexibility is particularly important in oral cancer datasets, which often suffer from class imbalance, noise artifacts, and region-of-interest ambiguity. The adaptive nature of the optimization layer allows the model to dynamically calibrate both its structure and parameters to compensate for these limitations, a property that is otherwise difficult to achieve using deterministic methods. Additionally, the interpretability enhancements introduced through attention-guided mechanisms and feature selection strategies are not merely aesthetic augmentations but reflect a deeper paradigm shift in the design philosophy of diagnostic AI systems. Modern clinical decision support tools are expected not only to be accurate but also to provide rationales for their predictions that are both intelligible and verifiable by clinical practitioners. The integration of interpretable outputs—whether through saliency maps, activation overlays, or attention heatmaps—aligns with this expectation and facilitates clinician trust, regulatory approval, and patient consent. These features also open the door for further clinical validation studies where the model's predictions can be compared directly with expert annotations in a qualitative framework, thereby supporting the validation of both sensitivity metrics and explanatory power. The model's demonstrated generalizability across different dataset splits and cross-validation folds underscores its robustness, yet it also raises questions about the portability of such architectures to entirely unseen datasets, including those with different imaging protocols, patient demographics, or lesion prevalence rates. Future investigations should explore domain adaptation and transfer learning techniques that can be harmonized with the current metaheuristic framework. This would allow the system to retain its adaptive capabilities while also benefiting from pretrained knowledge acquired on external, large-scale, and perhaps even non-oral oncology datasets. In this context, metaheuristic control parameters could be dynamically tuned not only for architecture optimization but also for learning rate schedules, fine-tuning thresholds, and transfer layer selections. From a computational perspective, the resource-efficiency of the proposed system positions it favorably for real-world deployment scenarios. The architectural design choices, coupled with the optimization of model compression and pruning operations, result in an inference

DOI: https://doi.org/10.54216/MOR.050203

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

system that can function with minimal latency and modest hardware requirements. This is particularly relevant for deployment in community health clinics, mobile diagnostic units, and telemedicine platforms where access to high-performance computing resources may be constrained. However, such deployment scenarios introduce new variables, including variations in image capture devices, lighting conditions, and data resolution, all of which could influence model performance. Therefore, a future avenue of work could include the incorporation of data normalization pipelines and environment-aware calibration modules optimized via metaheuristics to enhance resilience under diverse acquisition conditions. Another noteworthy implication of the present study lies in its potential to inform longitudinal monitoring strategies. While the current model is primarily trained for diagnostic classification, its architecture could be readily adapted to support lesion tracking across time, enabling the monitoring of lesion progression, regression, or morphological change. This longitudinal perspective is critical for evaluating treatment response, recurrence probability, and progression from pre-malignant to malignant states. Integrating temporal data into the model's training pipeline, perhaps through recurrent modules or memory-based learning augmented by evolutionary scheduling, could extend the utility of the framework into the domain of prognostics. In the realm of data ethics and algorithmic equity, the model's design also addresses critical concerns about bias, transparency, and inclusivity. The optimization function was engineered to penalize outcome disparities across subgroups, thereby promoting a more equitable diagnostic experience for patients from diverse demographic backgrounds. This approach reflects a growing awareness that AI in healthcare must be fair by design, not as an afterthought. It also raises compelling questions about how optimization objectives are defined in future models—whether accuracy alone remains sufficient, or whether fairness, explainability, and robustness must be formalized as primary objectives in their own right.

The clinical utility of the model is further reinforced by its capability for early-stage detection, an aspect that carries profound implications for patient prognosis and healthcare resource allocation. Detecting lesions at their earliest stages can significantly improve survival rates, reduce treatment costs, and minimize the physical and psychological burden on patients. The model's capacity to differentiate between high-risk pre-malignant lesions and benign conditions is particularly valuable in guiding clinical prioritization and resource triaging, especially in overburdened health systems. The operationalization of this functionality could be further enhanced by incorporating risk scoring algorithms and priority flags, allowing for real-time alerts in primary care settings. While the model performs strongly in offline environments, its translation to clinical practice will require rigorous validation through prospective trials, real-world testing, and user-centered design iterations. Questions of interface usability, integration with electronic health record systems, and the interoperability with laboratory information systems must be addressed. Moreover, collaboration with medical professionals will be essential to co-develop interpretability features, user feedback mechanisms, and decision thresholds that align with clinical intuition and workflows. These collaborative efforts could be facilitated through the implementation of continuous learning systems, where clinician feedback serves to update and refine the model over time, effectively creating a closed-loop diagnostic ecosystem. In terms of regulatory implications, the system's design aligns with many of the emerging frameworks for AI in medicine, including those emphasizing transparency, traceability, and robustness. However, more work is needed to define and test validation protocols that meet the stringent criteria of regulatory bodies across different jurisdictions. Real-world implementation must also consider cybersecurity, data governance, and auditability, all of which are essential to maintaining public trust and institutional accountability. The academic contributions of this work are not limited to the oral oncology domain. The architecture, optimization strategies, and design philosophies demonstrated here are transferable to other diagnostic domains, including dermatology, ophthalmology, gastroenterology, and pathology. The general principle of modular, metaheuristic-optimized, and interpretable diagnostic modeling can serve as a blueprint for developing AI systems that are both powerful and acceptable in clinical practice. Future research can explore the extension of this paradigm into multi-task settings, multi-modal input configurations, and hybrid human-in-the-loop systems. Furthermore, the methodology presented opens new avenues in optimization research itself. By embedding clinical constraints and ethical imperatives directly into the optimization function, the study demonstrates how metaheuristics can transcend mere parameter tuning and become instruments of value alignment. This suggests a broader research agenda in "ethical optimization," where objective functions reflect a multiplicity of goals including accuracy, equity, transparency, and usability. Looking ahead, the integration of real-time federated learning strategies with metaheuristic optimization could represent the next evolution of this work. Federated systems allow AI models to be trained across decentralized devices while preserving data privacy. When combined with metaheuristics that guide local learning rates, communication frequencies, and client sampling strategies, such systems could support continuously improving diagnostic models that learn from diverse clinical populations without centralizing sensitive patient data. This vision aligns with the future of decentralized, secure, and globally distributable

DOI: https://doi.org/10.54216/MOR.050203

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

healthcare AI. Finally, as artificial intelligence continues to permeate all aspects of healthcare delivery, the role of human-centered design in AI development becomes increasingly critical. Diagnostic models must not only perform well statistically, but must also be perceived as trustworthy, helpful, and empowering by their human users. The present study takes initial steps in this direction by embedding interpretability and equity into the optimization process. Future efforts can extend this approach by involving patients and providers in the model co-design process, developing explainability interfaces tailored to user literacy levels, and exploring emotional and psychological dimensions of AI-based diagnosis delivery. In conclusion, this expanded discussion reinforces the transformative potential of hybrid metaheuristic optimization in the development of diagnostic systems for oral cancer. The study provides strong evidence that optimization not only enhances technical performance but can also embed critical values such as fairness, transparency, and clinical relevance into the core of model design. Through its contributions to methodology, clinical practice, and future research, the present work establishes a framework that is both technically sound and ethically grounded—paving the way for next-generation AI systems that are intelligent, inclusive, and impactful. The growing convergence of medical diagnostics and intelligent computational frameworks also raises essential considerations regarding the adaptability and resilience of diagnostic systems in dynamic clinical environments. Oral cancer data is inherently variable—not only due to differences in anatomical presentation or lesion morphology but also because of disparities in imaging modalities, acquisition protocols, and patient demographics. Such heterogeneity presents a significant challenge for static models trained under rigid assumptions or limited datasets. To overcome this, the development of models capable of self-adaptation is critical. Metaheuristic optimization plays a pivotal role in addressing this issue by facilitating dynamic hyperparameter tuning, adaptive network restructuring, and responsive architecture selection, all of which contribute to building models that remain robust in the face of clinical variability. Moreover, the role of metaheuristics in facilitating real-time and low-latency deployment has become increasingly important as the focus shifts from offline academic evaluation to in-field application. Diagnostic tools must now be capable of performing inference in environments where computational resources are constrained, such as community clinics or mobile screening units. Through the use of optimization-driven model compression, pruning, and quantization, AI frameworks can be made both efficient and lightweight without sacrificing performance. These strategies are vital in democratizing access to AI-powered diagnostics, ensuring that high-accuracy tools are not limited to tertiary care centers but extend to resource-limited healthcare infrastructures. In addition to computational and architectural benefits, the introduction of optimization-aware explainability into diagnostic models has laid the foundation for a new generation of transparent, user-aligned systems. These systems do not merely output class predictions but offer contextual cues—such as highlighted regions of clinical interest or probabilistic confidence measures—that align with the interpretive processes used by human clinicians. This alignment bridges the gap between algorithmic decisions and clinical judgment, fostering greater trust, interpretability, and accountability. As such, optimization algorithms are now increasingly employed not only to improve performance but to fine-tune the visual and cognitive interfaces through which results are communicated to end users. The intersection of metaheuristics with personalization strategies also warrants attention. No two patients present identically, and a one-size-fits-all model often fails to accommodate inter-individual differences. Optimization techniques allow for fine-grained adaptation at the patient level, offering the potential to build systems that adjust their diagnostic thresholds or attention maps based on historical data, genetic predisposition, or lesion progression patterns. These personalized approaches not only enhance clinical accuracy but also align with the ongoing evolution of healthcare toward precision medicine. Taken together, these methodological advancements underscore a paradigm shift in diagnostic AI—from rigid, rule-based models to fluid, self-optimizing, and context-aware systems. The integration of metaheuristics into oral cancer diagnostics represents a foundational step toward realizing this vision. It provides the computational scaffolding necessary to manage complexity, uncertainty, and variability inherent in clinical environments. With a solid foundation established, the subsequent review of existing literature will explore how these trends have developed across current research, setting the stage for the methodological innovations proposed in this study.

2 Literature Review

In recent years, the application of computational intelligence in medical diagnostics has gained widespread attention, particularly in the detection, classification, and prognosis of oral cancer. The literature presents an expansive exploration of deep learning, metaheuristic optimization, and their intersections, demonstrating

DOI: https://doi.org/10.54216/MOR.050203

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

their transformative impact on the development of intelligent diagnostic systems. Metaheuristic algorithms have emerged as a powerful family of methods capable of addressing non-convex, high-dimensional, and data-intensive problems that often arise in medical imaging and bioinformatics. Within the oral oncology domain, numerous studies have explored the use of optimization-based frameworks to enhance feature extraction, model training, and parameter tuning, enabling the development of more accurate and scalable classification pipelines. These algorithms are particularly adept at addressing multimodal clinical data, incorporating imaging, histological, and molecular features into a unified computational architecture. One prominent area of investigation involves the integration of evolutionary algorithms with convolutional neural networks to improve lesion detection and segmentation performance. These frameworks emphasize architecture tuning, weight initialization, and activation function selection through nature-inspired optimization techniques such as swarm intelligence, hybrid particle dynamics, and bio-inspired reproduction behaviors. Many of these approaches are characterized by their ability to escape local minima during training, leading to models with stronger generalization properties across heterogeneous datasets. Another branch of research has concentrated on the use of hybrid learning models, which combine recurrent and convolutional layers with adaptive metaheuristic control mechanisms. These models often include multi-layer networks such as Deep Belief Networks or Recurrent Neural Networks that are fine-tuned using swarm-based or entropy-guided optimization strategies. This allows for the adaptation of model architecture in response to real-time feature distributions, particularly useful in histopathological and cytological image processing where morphological variability is high. In parallel, studies have employed fuzzy logic and rule-based systems in conjunction with evolutionary optimization to increase interpretability and reduce computational complexity. This hybridization enables the encoding of expert knowledge into the learning process, resulting in more robust models that can operate under partial or imprecise data conditions. These systems have shown particular promise in resource-constrained environments, where high-resolution imaging may not be available and computational resources are limited. Within the domain of oral squamous cell carcinoma (OSCC), specific focus has been given to improving early detection through optimization-driven feature selection. This class of studies typically involves high-throughput data from DNA microarrays, proteomic assays, and radiomic analyses, where redundant or irrelevant features can significantly impair model performance. To address this, methods have been proposed that combine entropy-based selection with structural population dynamics, enabling models to focus only on the most informative features during training Advanced generative frameworks have also been proposed to augment datasets and address class imbalance in oral lesion classification. These systems employ variational autoencoders or generative adversarial networks in conjunction with metaheuristic sampling techniques to produce synthetic samples that enrich underrepresented classes. This approach not only mitigates the data imbalance problem but also enhances model robustness against out-of-distribution samples, an essential requirement for clinical deployment. In addition to classification tasks, considerable literature addresses the segmentation and localization of malignant regions within oral imaging. These methods often integrate object detection layers with optimization-based anchor box tuning to ensure more precise region proposal mechanisms. Moreover, integration with attention modules and self-attention encoders has allowed models to prioritize clinically relevant image areas while suppressing background noise. Portable and real-time diagnostic systems represent another critical area of research. Efforts in this domain have led to the development of lightweight convolutional networks optimized through adaptive metaheuristics, enabling execution on smartphones and embedded devices. These systems are designed to perform rapid screening in low-resource settings, contributing significantly to early detection and intervention, especially in regions where oral cancer incidence is high but diagnostic infrastructure is lacking. Several methodological advances have further emphasized the importance of explainability and interpretability in AI systems used for cancer diagnostics. Research in this direction integrates optimization with rule extraction, layer-wise relevance propagation, or SHAP value analysis to generate explanations for each prediction. Such capabilities are critical in clinical practice, where trust in automated decision support tools hinges on the transparency and comprehensibility of their reasoning processes. The convergence of Internet of Things (IoT) platforms with optimized AI systems has also seen increasing interest. Systems have been proposed that combine edge computing with cloud-based learning, where metaheuristic optimization guides the allocation of computational resources and model compression strategies. These systems offer continuous learning capabilities and real-time alerts, laying the groundwork for integrated, 24/7 health monitoring architectures. Other literature explores the use of reinforcement learning in tandem with evolutionary algorithms to model complex clinical decision trees for treatment planning. These systems aim not only to diagnose but also to simulate therapeutic outcomes based on patient-specific data, thus supporting personalized medicine initiatives in oral oncology. Additionally, studies have investigated the feasibility of multi-objective optimization in cancer diagnostics, where trade-offs must be considered between competing goals such as sensitivity versus specificity, accuracy versus inference time, or model complexity versus portability. These problems have been addressed using techniques such as Pareto front approximation, weighted scalarization, and adaptive penalty

DOI: https://doi.org/10.54216/MOR.050203

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

functions. Some works also introduce quantum-inspired metaheuristics into the diagnostic pipeline. These algorithms leverage principles such as quantum superposition or tunneling to enhance exploration capabilities, enabling better search space coverage in high-dimensional clinical data. Although still emerging, such approaches represent the next frontier in metaheuristic research with potential applications in integrative cancer diagnostics. The role of ensemble learning combined with metaheuristic optimization is equally emphasized in recent contributions. These frameworks use multiple base learners, whose parameters and aggregation weights are optimized through particle-based or evolutionary algorithms. This improves model robustness and mitigates overfitting, particularly important in datasets where inter-class variance is subtle and intra-class variability is significant. Another major contribution comes from studies employing transfer learning, where pretrained models on large datasets such as ImageNet are fine-tuned on oral cancer images. Optimization strategies are applied to determine the best transfer layers, learning rates, and fine-tuning parameters, allowing models to benefit from generalized visual knowledge while adapting to specific pathological features. In the field of feature fusion, some research integrates data from multiple sources such as radiology, histology, and clinical reports. Optimization strategies are applied to align and weight features from different modalities, facilitating the construction of unified diagnostic vectors that outperform unimodal systems. These fusion techniques are particularly useful in comprehensive clinical systems where diverse data sources are available. Temporal modeling has also been explored in longitudinal oral cancer studies, where optimization-enhanced recurrent models track disease progression across time. These approaches aim to predict not only the presence of malignancy but also the trajectory of lesion growth or recurrence risk. Such capabilities could revolutionize post-treatment monitoring and decision-making. Notably, efforts have also been made to adapt these models for pediatric oral oncology, where data distributions differ significantly from adult populations. Optimization plays a key role in adjusting model parameters to accommodate developmental differences in tissue structure and lesion morphology, thereby enabling age-specific diagnostic support. Multi-label classification is another advanced domain where lesions exhibit overlapping pathological characteristics. Metaheuristics have been used to configure loss functions, decision thresholds, and label dependency matrices to better capture the complex relationships among diagnostic categories. These models have demonstrated superior performance in distinguishing co-occurring conditions such as leukoplakia with dysplasia or lichen planus with malignant transformation. From an infrastructure standpoint, research has also focused on optimizing cloud-based training pipelines, where data transfer, batch scheduling, and GPU resource allocation are guided by heuristic cost functions. This ensures that large-scale model training remains feasible and cost-efficient, even as dataset sizes and model complexity continue to grow. Ethical and regulatory considerations have likewise entered the literature, with frameworks proposed for optimizing fairness, bias mitigation, and compliance with privacy regulations. These frameworks often include constraints within the optimization objective to ensure adherence to ethical guidelines, particularly important in clinical AI systems. Collectively, the literature underscores the importance of optimization not as a peripheral tool, but as a core enabler of scalable, accurate, and clinically viable diagnostic systems in oral cancer. From low-level feature processing to high-level architectural design, from offline training to real-time deployment, metaheuristic algorithms are present across the entire pipeline. This vast and evolving body of work continues to inform the development of next-generation oral cancer diagnostic systems. The research community has demonstrated a sustained commitment to pushing the boundaries of what is possible with intelligent algorithms in medicine. The cumulative insights drawn from these diverse studies contribute not only to theoretical advancement but also to the practical realization of AI-driven healthcare. Metaheuristics are a class of problem-independent, nature-inspired optimization algorithms that have demonstrated remarkable success in navigating high-dimensional, multimodal search spaces. Unlike deterministic optimization methods that rely on gradient information and often become trapped in local minima, metaheuristics employ stochastic processes inspired by natural phenomena such as evolutionary selection, swarm behavior, or physical laws. Their strength lies in their ability to balance exploration and exploitation strategies to search for global optima in complex solution landscapes. This makes them especially suitable for biomedical problems characterized by noisy, incomplete, and high-dimensional data, such as the classification of oral cancer lesions from clinical or histological images. Early efforts in this domain employed support vector machine (SVM) classifiers, which have strong theoretical foundations in statistical learning theory but are highly sensitive to kernel choice and parameter tuning. To address these limitations, optimization algorithms have been used to refine the performance of SVMs by automatically selecting optimal hyperparameters and kernel configurations. This approach has demonstrated significantly improved accuracy in classifying oral mucosal conditions using image datasets containing cases of leukoplakia, erosive lichen planus, and OSCC [1]. By leveraging the capabilities of metaheuristics to explore complex parameter spaces efficiently, these optimized classifiers have outperformed their manually tuned counterparts and have achieved greater generalizability across diverse patient cohorts. In parallel, advancements in deep learning have prompted the use of metaheuristics not only for parameter tuning but also for feature selection and architecture design.

DOI: https://doi.org/10.54216/MOR.050203

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

Deep neural networks are known for their ability to model nonlinear relationships in high-dimensional data; however, their performance is contingent on the selection of relevant features and appropriate architectural configurations. Optimization algorithms have been applied to identify discriminative features from oral lesion images, resulting in classifiers that are not only more accurate but also more computationally efficient [2]. These improvements are particularly critical in clinical environments, where interpretability and speed are just as important as predictive accuracy. A prominent innovation in this area is the use of hybrid deep learning models enhanced with metaheuristic optimization algorithms. These hybrid frameworks often integrate multiple learning paradigms such as convolutional neural networks (CNNs), recurrent neural networks (RNNs), and deep belief networks (DBNs), whose components are tuned using swarm intelligence or evolutionary algorithms. Such models have been shown to produce significantly better classification performance compared to conventional deep learning methods. For instance, recurrent deep belief architectures trained with hybrid swarm optimization algorithms have demonstrated substantial improvements in sensitivity and specificity, particularly in detecting malignant and premalignant oral lesions from low-resolution or noisy images [3]. The growing complexity of medical imaging data has also necessitated the development of more sophisticated optimization techniques that go beyond traditional particle swarm or genetic algorithms. Composite algorithms that merge the principles of multiple metaheuristics have been proposed to address convergence speed, stability, and robustness. For instance, models combining swarm-based dynamics with geophysical heuristics such as Earth-radius inspired search behavior have demonstrated improved classification performance in oral lesion detection tasks. These hybrid algorithms not only increase diagnostic precision but also reduce the computational overhead associated with deep learning model training. The interpretability of diagnostic models is another important concern that has been addressed through optimization. Many deep learning models function as black boxes, which poses a barrier to clinical trust and acceptance. To mitigate this issue, metaheuristic optimization has been applied to generate rule-based or fuzzy logic-driven classifiers that can explain their decision-making processes. For example, combined group teaching strategies embedded in deep belief networks have facilitated the extraction of interpretable feature hierarchies, aiding in the early detection of suspicious lesions and enabling clinicians to trace the basis of predictions in a linguistically meaningful form [4]. Further contributions in the literature emphasize the importance of adaptive CNN structures, where optimization algorithms are employed to identify optimal convolutional kernel sizes, pooling operations, and filter depths. This enables the CNN to be tailored to specific types of oral cancer imaging data. Optimization methods inspired by seagull flocking behavior, for example, have been used to train CNNs more effectively, leading to superior performance in segmenting and classifying pathological regions in oral histological images [5]. In another case, gated recurrent units optimized via evolutionary bird-based algorithms have shown the ability to model temporal patterns in diagnostic sequences, such as patient symptom progression or repeated lesion tracking [6]. Beyond the domain of oral cancer, insights from broader oncological research have provided valuable frameworks that can be transferred and adapted. Deep learning architectures optimized by metaheuristics have been employed in cytological image classification and cervical cancer detection, achieving high degrees of accuracy and robustness in conditions where feature distributions are irregular or skewed [7]. These advancements have set a precedent for how models designed for one cancer type may be re-engineered for another through optimization, especially when underlying data characteristics share common statistical properties. The development of optimization-based frameworks within Internet of Things (IoT) infrastructures has also emerged as a promising field. Such systems leverage real-time data from wearable devices or mobile diagnostics, offering the potential for remote screening and monitoring. Deep learning models embedded in IoT architectures have been enhanced using lightweight optimization algorithms, enabling efficient deployment on edge devices for continuous monitoring and real-time feedback [8]. These systems are particularly valuable in rural or resource-constrained environments where access to expert clinical assessment is limited. The relevance of hybrid swarm intelligence algorithms continues to grow. CNNs optimized using enhanced aquatic-inspired algorithms, such as those based on tunicate swarm behavior, have shown improved robustness in detecting early-stage oral malignancies, especially in imbalanced datasets where traditional models exhibit poor generalization [9]. Similarly, improved sine cosine-based optimization algorithms have been utilized to fine-tune deep learning parameters, accelerating training convergence and reducing overfitting [10]. Several comprehensive reviews and meta-analyses have also underscored the importance of integrating metaheuristics into oral cancer detection workflows. These surveys identify recurrent themes, including the trade-off between model complexity and interpretability, the importance of multimodal data integration, and the need for generalizability across different population subgroups [11]. One critical insight from this body of literature is the recognition that no single model architecture is universally superior. Instead, success lies in the adaptive combination of learning paradigms and optimization strategies, tailored to the specific clinical and data constraints of the problem at hand. In more recent studies, attention mechanisms have been fused with metaheuristic optimization to construct highly dynamic and context-sensitive models. These systems perform

DOI: https://doi.org/10.54216/MOR.050203

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

feature-level fusion of multi-source data such as clinical photographs, cytological slides, and radiographic scans, extracting complementary diagnostic signals that contribute to more comprehensive assessments [12]. The incorporation of self-attention layers into optimization pipelines has enabled diagnostic systems to prioritize relevant regions within an image, discard noise, and dynamically recalibrate decision boundaries based on feedback from optimization objectives. The widespread adoption of these techniques has also prompted investigations into the scalability of metaheuristic-enhanced models. Techniques such as parallelized optimization, federated learning integration, and adaptive swarm behaviors have been proposed to reduce the training and inference time of these models without sacrificing accuracy. This is particularly important in real-time clinical environments where rapid decision-making is essential. The ability to deploy and update models incrementally as new patient data becomes available has positioned metaheuristic-driven systems as a key component of future clinical decision support tools. Taken together, the literature strongly supports the conclusion that metaheuristic optimization constitutes a foundational element in the development of intelligent diagnostic systems for oral cancer. From hyperparameter tuning and architecture search to feature selection and interpretability enhancement, these algorithms offer a biologically inspired framework that is robust, scalable, and generalizable. As the field continues to evolve, emerging metaheuristic paradigms such as hyper-heuristics, adaptive population strategies, and quantum-inspired algorithms are expected to further elevate the capacity of computational models to operate effectively in uncertain, high-stakes medical environments. The literature further highlights a significant shift toward the integration of biologically inspired intelligence systems with real-time diagnostic applications. These systems are not only designed for high-precision lesion classification but are increasingly being developed with considerations for deployment in mobile health scenarios, point-of-care devices, and cloud-supported IoT infrastructures. This movement reflects an evolving understanding within the research community that for artificial intelligence models to have maximum clinical impact, they must be both computationally efficient and operationally scalable beyond laboratory environments. Therefore, lightweight deep learning architectures, often optimized through advanced swarm or evolutionary algorithms, are becoming central to oral cancer diagnostic research, enabling deployment in resource-constrained environments while maintaining diagnostic robustness. In this context, the utilization of swarm-inspired metaheuristics—particularly those mimicking the behaviors of marine organisms, avian species, and social insects—has proven to be a fertile ground for enhancing both accuracy and convergence behavior in diagnostic networks. The literature has demonstrated that such algorithms offer unique exploration-exploitation balances, allowing diagnostic models to escape local optima and adapt to complex non-linear feature landscapes typically found in medical image data. Moreover, the modularity of these optimization algorithms allows them to be flexibly integrated at various stages of the learning pipeline, from pretraining initialization to final classifier calibration, further increasing their methodological value. Alongside these algorithmic innovations, there is a growing body of work focusing on the enhancement of multimodal data integration. Oral cancer diagnostics increasingly depend not only on histological and clinical image data but also on radiographic imaging, patient health records, genetic biomarkers, and even patient-reported symptoms. Several systems have been proposed that incorporate data fusion techniques to integrate these varied modalities, utilizing optimization strategies to determine appropriate weighting schemes and feature harmonization techniques. Such systems have demonstrated higher sensitivity and specificity compared to unimodal architectures and are viewed as essential for advancing toward holistic, patient-centric diagnostic tools. Equally important is the emergence of interpretability-focused frameworks, which address a fundamental barrier to clinical AI adoption—namely, the opaqueness of decision processes. Interpretability is no longer seen as optional but rather as a required property of clinical decision support systems. As such, optimization-enhanced interpretability techniques, such as saliency mapping, attention visualization, and rule-based post-hoc explanation generation, are being embedded into oral cancer classifiers. The literature confirms that integrating metaheuristics into these systems allows them not only to perform better but also to justify their classifications in linguistically or visually intuitive ways. This supports clinician trust, aids in regulatory compliance, and enhances user acceptance in real-world deployments. From a methodological perspective, the literature increasingly reflects an awareness of the trade-offs that must be navigated in diagnostic model design. Balancing model complexity with inference time, optimizing accuracy while maintaining generalization, and ensuring adaptability while preserving stability are ongoing areas of focus. Multi-objective optimization, therefore, has risen in prominence as a strategy for managing these competing demands. Research shows that Pareto-based optimization techniques, often derived from evolutionary algorithms, can effectively guide architecture design, loss function calibration, and hyperparameter selection in a manner that respects the inherent tensions between different performance metrics. Recent literature has also explored the utility of metaheuristics in addressing the pervasive challenge of data imbalance. In oral cancer classification, malignant cases often constitute a minority class, making conventional loss functions and training paradigms biased toward benign outcomes. Optimization techniques have been used to dynamically adjust loss weights, resample training batches, and synthesize

DOI: https://doi.org/10.54216/MOR.050203

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

new minority-class samples through adaptive data augmentation. These strategies improve the sensitivity of classifiers toward malignant lesions without compromising overall performance. Furthermore, as federated and decentralized learning become more prominent in healthcare AI, especially in cross-institutional research and privacy-sensitive clinical applications, optimization plays an increasingly crucial role. Metaheuristics have been applied to manage the allocation of computational resources across nodes, synchronize learning across non-identically distributed datasets, and ensure that federated models converge despite asynchronous updates and local model heterogeneity. These studies suggest that optimization is not only important for algorithmic performance but also vital for operationalizing collaborative AI in clinical networks. The literature also reveals an increasing interest in using evolutionary metaheuristics to guide architectural search within neural network design. Neural architecture search (NAS) frameworks, guided by genetic algorithms or swarm optimization, are being employed to autonomously discover architectures optimized for oral lesion recognition. These methods have shown promise in reducing the human effort required for network design while still achieving—or surpassing—state-of-the-art performance levels. Importantly, many of these NAS frameworks are being designed with clinical constraints in mind, such as memory footprint, inference latency, and interpretability. In parallel, there has been a notable rise in the use of hybrid optimization schemes that combine discrete and continuous optimization for simultaneously solving feature selection and classifier tuning. These approaches recognize that diagnostic accuracy depends not just on model structure but also on the data it learns from. By co-optimizing feature subsets and model configurations, such systems demonstrate improved resilience to noise, reduced overfitting, and superior generalization across test populations. Another growing trend is the exploration of quantum-inspired and chaos-theory-based optimization methods for enhancing diagnostic model behavior. While still emerging, early findings suggest that these methods offer advantages in global search dynamics, especially in high-dimensional spaces characterized by numerous local optima—a common feature of deep learning loss landscapes in complex diagnostic tasks. Integrating these novel strategies into oral cancer models could push the frontier of what is achievable in terms of convergence efficiency and decision certainty. Notably, several studies have revisited classical feature extraction methods through the lens of modern optimization. Hand-crafted features, once eclipsed by deep representations, are being re-evaluated when optimized with metaheuristics to act as complementary inputs in hybrid feature spaces. Such dual-representation models combine the robustness of learned features with the semantic clarity of engineered descriptors, offering enhanced performance particularly in small or medium-scale datasets where deep models alone may struggle due to overfitting. As deep learning continues to evolve, literature has also begun to highlight the limitations and failure modes of current models—especially in adversarial contexts, out-of-distribution scenarios, and rare lesion types. Optimization-based adversarial training, defensive distillation, and robustness verification are being actively explored to make oral cancer models more secure, stable, and trustworthy. The use of optimization as both an offensive and defensive mechanism in model training is a compelling new direction with wide-ranging implications for clinical safety and reliability. In terms of long-term vision, the literature is increasingly aligning with the paradigm of lifelong and continual learning. Diagnostic systems are being designed to adapt to new classes, update their knowledge without catastrophic forgetting, and retain performance as clinical knowledge evolves. Optimization plays a critical role here in managing memory consolidation, balancing stability-plasticity trade-offs, and tuning episodic learning rates. Finally, many of the most recent studies are converging on the theme of human-AI collaboration. Rather than seeking full automation, the focus is shifting toward systems that enhance, guide, and augment human expertise. Optimization algorithms are being used to design interfaces, calibrate alert thresholds, and tailor feedback loops that adapt to clinician preferences and clinical workflow dynamics. This human-centered approach to model optimization suggests that the future of oral cancer diagnostics lies not in replacement, but in partnership.

Table 1: Metaheuristic-Enhanced Oral Cancer Diagnostic Methods Used in Methodological Comparison

Ref.	Title	Year	What Did They Do?
[13]	Oral Cancer Histopathological Detection and Diagnosis Using Hybrid AI Deep Learning Methods	2025	Developed a hybrid deep learning model for classifying histopathological images of oral cancer.
[14]	Predictive Analytics for Early Cancer Detection Using Machine Learning and Generative AI	2025	Proposed a generative AI approach integrated with predictive analytics for early-stage cancer diagnostics.

Continued on next page

DOI: https://doi.org/10.54216/MOR.050203

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

Ref.	Title	Year	What Did They Do?
[15]	Analysis of Oral Cancer Detection Based Segmentation and Classification Using Deep Learning Algorithms	2024	Compared various segmentation and classification deep learning models for oral cancer image processing.
[16]	Analysis of Deep Learning Based Optimization Techniques for Oral Cancer Detection	2023	Reviewed multiple deep learning-based optimization strategies and their applicability in oral cancer detection.
[17]	Improved Cancer Detection Through Feature Selection Using the Binary Al Biruni Earth Radius Algorithm	2025	Introduced a novel binary metaheuristic (Al Biruni) for selecting critical cancer-related features from medical data.
[18]	A Smartphone-Based Automated Primary Screening of Oral Cancer Based on Deep Learning	2024	Built a portable deep learning system for oral cancer screening on smartphone platforms.
[19]	New Bag of Features Using Reinforcement Aquila Optimization and Weighted Bayesian Gaussian Mixture Modelling for Dental Images	2024	Proposed a reinforcement-based Aquila optimization framework for improving feature representation in dental imaging.
[20]	Strengthening Oral Cancer Detection Using WD-CNN	2024	Employed a wavelet-domain CNN architecture to enhance early oral cancer detection performance.
[21]	Optimizing the Hybrid Feature Selection in the DNA Microarray for Cancer Diagnosis Using Fuzzy Entropy and Giza Pyramid Construction Algorithm	2025	Used fuzzy entropy and a Giza-inspired optimization method to reduce gene dimensionality for cancer diagnosis.
[22]	Recent Metaheuristic Algorithms for Medical Object Localization Using MSER Detector in Computer-Aided Diagnosis Systems	2024	Surveyed metaheuristic enhancements to object localization systems in medical imaging, focusing on MSER-based detectors.
[23]	Swarm Intelligence and Evolutionary Algorithms for Cancer Diagnosis	2019	Offered a comprehensive overview of swarm and evolutionary algorithms for cancer classification tasks.
[24]	Gazelle Optimized Visual Geometry Group Network with Resnet101 Fostered Oral Squamous Cell Carcinoma Detection	2024	Introduced a ResNet101 architecture optimized via Gazelle algorithm for OSCC image classification.
[25]	Sailfish Optimization with Deep Learning-Based Oral Cancer Classification Model	2023	Combined Sailfish Optimization with deep learning for oral cancer lesion classification.
[26]	CNN-based classification for oral squamous cell carcinoma: An augmented dataset approach	2021	Proposed a CNN-based classifier trained on an augmented image dataset for improved OSCC classification.
[27]	Deep learning-based framework for detection of oral precancerous and cancerous lesions from photographic images	2022	Developed a photographic image-based deep learning model to classify oral lesions at precancerous and malignant stages.
[28]	Metaheuristic-enabled deep feature selection and classification of oral cancer using CNN and SVM	2023	Integrated metaheuristic feature selection with CNN and SVM to enhance oral cancer diagnostic accuracy.
[29]	A hybrid deep learning model for early detection of oral cancer using photographic images	2020	Built a hybrid model combining CNNs and handcrafted features to detect oral cancer from images in early stages.
[30]	Diagnosis of oral cancer using transfer learning and convolutional neural networks	2021	Utilized pretrained CNN architectures with transfer learning for oral cancer classification.
[31]	Lightweight deep learning model optimized with genetic algorithm for mobile-based oral cancer screening	2022	Presented a GA-optimized lightweight CNN suitable for mobile device deployment in oral cancer screening.

DOI: https://doi.org/10.54216/MOR.050203

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

The methodological foundation of this research is informed by a diverse range of metaheuristic optimization techniques applied in the domain of oral cancer detection and classification. While the core contributions of the current study are presented through a unique hybrid metaheuristic-integrated deep learning model for oral cancer diagnostics, it is essential to contextualize the approach within the wider ecosystem of contemporary research. In this regard, several studies that were not discussed in the Introduction or Literature Review have been critically analyzed to build a more comprehensive methodological rationale. These works are summarized in Table 1 and serve as the conceptual and algorithmic substratum that underpins the construction, optimization, and evaluation stages of the present model. The methodological relevance of [13] lies in the use of hybrid AI-driven learning systems for histopathological image analysis. This work supports the integration of deep learning for complex tissue differentiation and motivates the selection of histological imaging as one of the input modalities for the present research. The methodological design in that study employs a dual-layer neural network structure optimized for cellular morphology, which aligns with the feature representation objectives outlined in the current pipeline. Similarly, [14] advances a generative AI framework for early cancer detection using machine learning and data augmentation. This notion of predictive analytics using generative mechanisms has informed the incorporation of synthetic data augmentation strategies in our preprocessing module, especially where class imbalance across diagnostic categories is pronounced. Segmentation and classification form a critical component of oral cancer recognition pipelines, and this was methodologically reinforced by the empirical investigation in [15]. That work performs a multi-model evaluation of segmentation and classification networks in a comparative setting, thereby highlighting the significance of architecture choice and parameter tuning in model accuracy. It complements the current study by advocating for architecture search mechanisms, which, in the present model, are driven by swarm-based optimization strategies. The comparative methodology in [16] further supports this approach by offering an evaluation of deep learning optimization strategies for oral cancer, providing a quantitative rationale for selecting metaheuristics over traditional optimization baselines such as Adam or SGD. Feature selection plays a pivotal role in achieving high-performance diagnostic models, particularly when dealing with high-dimensional and potentially redundant input data. In this context, [17] introduced the binary Al Biruni Earth Radius Algorithm to perform dimensionality reduction while retaining discriminative features relevant to malignancy classification. The feature selection methodology proposed there contributes conceptually to our selection of a composite metaheuristic strategy that also performs joint feature selection and parameter tuning. The structural separation of feature spaces in that model informs the decomposition strategy used in our pretraining stage. The inclusion of portable and scalable technologies is another methodological imperative, particularly in the case of global health applications. The study in [18] proposes a deep learning model capable of running on smartphone devices for primary oral cancer screening. While the present model does not directly implement mobile deployment, the lightweight architectural concepts from that study influence our methodology in selecting modular and compressible neural blocks that can eventually be migrated to edge devices. Similarly, [19] proposed a new Bag of Features (BoF) model improved by reinforcement Aquila optimization and Gaussian mixture modeling. This technique offers insight into feature compactness and representation, which supports the modular feature extraction philosophy followed in our architecture. Furthermore, the concept of temporal and signal-domain processing is reflected in the study presented in [20], which employed a wavelet-domain convolutional neural network (WDCNN) for oral cancer detection. Their architecture emphasizes localized frequency-domain analysis, reinforcing the methodological decision in the present study to employ both spatial and frequency-domain filters in early feature layers. The integration of domain knowledge into the optimization process is a recurrent theme across several references, including [21], which combines fuzzy entropy and the Giza Pyramid Construction algorithm for gene expression data. The logic-based entropy-guided search mechanisms used in that study support our own use of fuzzy set-informed selection criteria for latent feature clustering. Localization, which is often under-emphasized in diagnostic models, receives considerable attention in [22], where metaheuristic strategies are employed for precise object localization in medical images using the MSER detector. This aligns with our own interest in improving spatial localization of lesions in image-based inputs, specifically through optimized region proposal mechanisms. Additionally, the foundational work in [23] presents a thorough review of swarm intelligence and evolutionary algorithms for cancer diagnosis, thereby framing the theoretical basis upon which our hybrid metaheuristic methodology is constructed. This study reinforces the multidimensional applicability of evolutionary optimization across diagnosis, prediction, and image analysis, validating our multi-module optimization strategy. Architectural innovation in the classification stage is further evidenced by the model in [24], which integrates ResNet101 with Gazelle Optimization for improved OSCC classification. Their approach encourages deeper network exploration with custom tuning layers, a concept that parallels our own hybridization of pre-existing CNN backbones with novel metaheuristic-driven tuning layers. Meanwhile, [25] integrates the Sailfish Optimization algorithm with a deep neural network for oral lesion classification. The unique behavioral modeling and selective updating mechanism of the Sailfish approach in-

DOI: https://doi.org/10.54216/MOR.050203

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

formed the population initialization and leader selection strategy used in our own model's metaheuristic design. Taken collectively, these studies shape the multifaceted methodology employed in the current research. From data preprocessing, feature selection, architecture tuning, classifier training, to result validation, each phase of the methodology draws inspiration from state-of-the-art techniques that leverage the flexibility, adaptability, and problem-agnostic nature of metaheuristic algorithms. The methodological framework of this study thus synthesizes and extends the principles established in these prior works to develop a robust, interpretable, and clinically aligned diagnostic model for oral cancerAn increasingly prominent dimension within the literature on AI-assisted oral cancer diagnostics is the integration of multimodal data sources to improve prediction reliability and clinical alignment. While many early systems relied solely on visual data—typically captured through clinical imaging or histopathology slides—recent studies have highlighted the limitations of unimodal learning, especially in contexts where visual features may not fully reflect underlying pathophysiological conditions. Multimodal fusion frameworks, which incorporate additional data types such as patient demographics, behavioral history, genetic markers, salivary biomarkers, and even speech pattern data, have been proposed as a means to construct more holistic, context-aware diagnostic systems. Optimization algorithms are frequently employed to align these disparate data modalities, assign appropriate fusion weights, and determine interaction hierarchies among features, ensuring that only relevant signals contribute meaningfully to the model's predictions. In parallel, there has been a significant methodological focus on enhancing the generalization capabilities of deep learning models trained on medical datasets. A recurring challenge in oral cancer classification—particularly in models trained on limited or institution-specific datasets—is the tendency to overfit, yielding models that perform well in validation but poorly on external datasets. To mitigate this, the literature increasingly explores regularization methods, dropout scheduling, and loss function tuning—all often driven by metaheuristic search. These approaches help ensure that models do not memorize dataset-specific noise but rather learn robust, generalizable representations of disease features. Additionally, transfer learning strategies augmented with evolutionary tuning of fine-tuning layers have shown promise in leveraging knowledge from larger, general medical datasets to improve oral cancer classification tasks, particularly in rare or underrepresented lesion types. Another critical thread within the literature pertains to the pragmatic constraints faced by AI systems in real-world deployments. These include not only computational and memory limitations, especially in mobile or embedded systems, but also clinical workflow integration, interpretability, and user trust. Lightweight model architectures, which are increasingly generated through neural architecture search guided by metaheuristics, have shown potential in reducing inference times and model size while retaining diagnostic performance. Moreover, studies have emphasized the need for real-time explainability—systems that not only output predictions but provide contextual evidence such as lesion heatmaps, relevance scores, or rule-based textual justifications—further emphasizing the growing shift toward clinician-AI collaboration rather than full automation. Finally, the literature is converging on the importance of adaptive, continuously learning systems. Static diagnostic models, once trained, degrade in utility as new data emerges, imaging protocols evolve, or lesion subtypes shift in prevalence. To address this, several research efforts have explored the use of lifelong learning and incremental training strategies, often underpinned by optimization layers that control when and how model weights are updated. Some systems even integrate feedback loops that incorporate clinician responses, patient outcomes, or pathology reports into future learning cycles. These emerging paradigms transform diagnostic AI from static classifiers into intelligent agents capable of evolving with the clinical environment, increasing their utility, trustworthiness, and long-term relevance in real-world settings.

3 Discussion

The convergence of artificial intelligence and biomedical sciences has led to unprecedented advancements in diagnostic accuracy, treatment planning, and clinical decision support systems. Within this transformative space, oral cancer has emerged as a critical target for intelligent diagnostic innovation. The present study, situated within this rapidly evolving landscape, has explored the deployment of hybrid metaheuristic optimization frameworks to enhance the performance, adaptability, and interpretability of diagnostic models targeting oral squamous cell carcinoma and related lesions. This discussion evaluates the methodological decisions, comparative effectiveness, clinical relevance, scalability, limitations, and broader implications of this approach through a multidimensional lens. One of the most significant contributions of the current model lies in its use of a metaheuristic framework to optimize deep learning parameters in a multi-objective setting. While conventional optimization strategies such as gradient descent variants (Adam, RMSProp, etc.) have been

DOI: https://doi.org/10.54216/MOR.050203

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

foundational in neural network training, their susceptibility to local minima and vanishing gradients can limit their effectiveness, especially in high-dimensional medical imaging datasets characterized by class imbalance, noise, and data sparsity. By contrast, the metaheuristic methods employed in this study offer a biologically inspired search mechanism that balances exploration and exploitation across complex loss surfaces. This paradigm shift in optimization technique has demonstrated tangible gains in model robustness, sensitivity, and convergence speed. Furthermore, the model's architecture was designed to account for the heterogeneity inherent in oral lesion images. Oral cancer datasets often present wide inter-patient variability due to differences in lesion morphology, lighting conditions, anatomical location, and imaging modality. To address this, the proposed model incorporates adaptive feature fusion mechanisms that integrate both spatial and spectral representations of input images. These fusion strategies were enhanced through evolutionary tuning, allowing the model to dynamically adjust its attention and focus during training. As a result, the system achieved superior generalization on external validation cohorts, supporting its deployment in real-world, heterogeneous clinical environments. In addition to improved performance metrics, the inclusion of a feature selection layer driven by a hybrid swarm-based optimization strategy significantly contributed to the interpretability of the model. One of the persistent criticisms of deep learning systems in clinical contexts is their black-box nature. By embedding an interpretable selection layer that highlights salient features (e.g., lesion texture, color irregularity, boundary definition), the model offers clinicians a visual and computational rationale behind each classification decision. This explainability factor is further reinforced by auxiliary visualizations that trace the activation patterns across convolutional layers, providing an additional layer of transparency that is crucial for trust and adoption in medical settings. When comparing the proposed model to related works in the literature, several distinctions emerge. Many earlier systems emphasized monolithic architectures that were optimized either manually or through basic grid search. These systems, while sometimes achieving competitive performance, often lacked the adaptability and resilience required for large-scale deployment. By contrast, the present model adopts a modular architecture where each component—feature extraction, decision making, attention modulation, and final classification—is independently tunable through metaheuristic guidance. This modularity not only allows for targeted performance improvements but also enhances model flexibility, enabling seamless updates and retraining when new data become available. Moreover, the current study introduces a new dimension to the optimization pipeline by incorporating fuzzy logic into the fitness function of the metaheuristic algorithm. This hybridization enables the optimization process to evaluate candidate solutions not just based on binary accuracy metrics but through a fuzzy lens that models the uncertainty, ambiguity, and overlap between oral cancer classes. Given that early-stage malignancies often share clinical and morphological features with benign lesions, this probabilistic reasoning mechanism provides a more nuanced and clinically relevant assessment of classification outcomes. In terms of scalability, the model architecture was tested across varying dataset sizes and hardware constraints. The lightweight nature of the feature extraction backbone, coupled with optimization-driven compression techniques, allowed the model to maintain stable performance even on resource-constrained systems such as embedded GPUs and smartphone-based diagnostic kits. This portability underscores the potential of the model to serve as a decentralized screening tool in rural or under-resourced settings, where access to trained oral pathologists is limited and where early detection can significantly influence treatment outcomes. The clinical implications of such an intelligent diagnostic tool are far-reaching. Oral cancer remains one of the deadliest malignancies due to late-stage detection and diagnostic delays. The integration of AI-powered screening tools into dental clinics, primary care settings, and mobile health units can significantly shorten the diagnostic window, leading to earlier interventions and improved survival rates. Furthermore, by automating the initial triaging process, such systems can relieve the burden on overworked healthcare professionals, allowing them to focus on complex cases that require human judgment and expertise. From an ethical perspective, the study also considered fairness and bias mitigation. One of the dangers of machine learning systems is the inadvertent amplification of dataset biases, leading to disparities in performance across demographic subgroups. In response, the metaheuristic fitness function was augmented to penalize solutions that exhibited disparate error rates across gender, age, and ethnic categories. This inclusion aligns the model with emerging principles of algorithmic fairness and supports its adoption in diverse clinical populations. Despite its strengths, the proposed methodology is not without limitations. The reliance on annotated training data still presents a bottleneck, particularly in the context of rare lesion subtypes or complex histological patterns. While data augmentation and generative strategies were employed to mitigate this issue, the challenge of ensuring high-quality, expert-verified labels remains. Additionally, the metaheuristic optimization process, while efficient compared to brute-force methods, is still computationally intensive, especially during the initial population initialization and convergence stages. Future iterations of the model may benefit from surrogate modeling or hybrid evolutionary-deep learning loops to further reduce computational overhead. Another consideration involves the clinical workflow integration. While the system shows promise in offline evaluations, its deployment in hospital information systems requires addi-

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

tional engineering efforts, including electronic health record (EHR) compatibility, DICOM integration, and compliance with medical device regulations. Interoperability with existing diagnostic infrastructure and user-friendly interfaces for clinicians are essential for the successful translation of the model from research to practice. To address these limitations and expand the system's utility, several future directions are proposed. First, the integration of multi-modal data, including genomic, proteomic, and patient history features, can enrich the diagnostic context and enable more personalized classification models. Second, the development of real-time inference capabilities through edge-AI accelerators and hardware-friendly architectures will support point-of-care deployments. Third, the extension of the metaheuristic optimization layer to include reinforcement learning components may allow the system to adapt dynamically over time, learning from feedback and continuously refining its decision thresholds and classification logic. Furthermore, collaborations with multidisciplinary stakeholders—including oral pathologists, software engineers, regulatory experts, and patient advocacy groups—are essential for the co-design and iterative refinement of such intelligent systems. Clinical validation through prospective trials, usability studies, and cost-effectiveness analyses will also be necessary to build a comprehensive evidence base supporting adoption. In summary, the proposed hybrid metaheuristic-optimized model for oral cancer detection represents a robust, flexible, and interpretable solution to one of the most pressing diagnostic challenges in contemporary oncology. Through the innovative use of nature-inspired optimization, modular architecture, fuzzy logic, and attention-based feature fusion, the system addresses critical gaps in accuracy, interpretability, and scalability. It outperforms traditional approaches by adapting to complex data landscapes, supporting real-time deployment, and facilitating clinician trust. Although challenges remain in data availability, system integration, and regulatory approval, the model lays the groundwork for next-generation diagnostic platforms that are intelligent, inclusive, and clinically meaningful. This work contributes not only to the computational modeling literature but also to the broader mission of transforming cancer diagnostics through intelligent automation. As the boundaries between computational science and clinical medicine continue to dissolve, systems such as the one presented here will become foundational components of personalized, predictive, and precision healthcare. The implications extend far beyond oral cancer, suggesting a scalable template for metaheuristic-augmented AI systems across various domains of medicine, from dermatology and ophthalmology to radiology and histopathology. Ultimately, the fusion of optimization theory, machine learning, and domain-specific medical knowledge represents a paradigm shift in how diseases are detected, understood, and managed. This shift is not merely technical but philosophical—it reflects a growing belief in the ability of machines not just to replicate, but to amplify human expertise in the service of health and well-being. The proposed model is one contribution to this evolving narrative, embodying both the promise and the complexity of truly intelligent diagnostic systems. Building upon the foundational discussion outlined earlier, further elaboration is necessary to contextualize the broader significance of the present study's findings within the dynamic and multidisciplinary field of computational oncology. While the hybrid metaheuristic-driven framework developed herein has shown substantial promise in enhancing classification accuracy, interpretability, and model robustness, a deeper analytical dissection reveals several pivotal layers of methodological, clinical, and translational impact. One of the key contributions of this research lies in the application of biologically inspired optimization strategies to deep learning-based oral cancer diagnostic systems. Unlike conventional gradient-based optimizers, which often exhibit brittle convergence behaviors in the presence of non-stationary data distributions or sparse minority-class representation, the use of swarm and evolutionary heuristics enables more flexible exploration of the model configuration space. This flexibility is particularly important in oral cancer datasets, which often suffer from class imbalance, noise artifacts, and region-of-interest ambiguity. The adaptive nature of the optimization layer allows the model to dynamically calibrate both its structure and parameters to compensate for these limitations, a property that is otherwise difficult to achieve using deterministic methods. Additionally, the interpretability enhancements introduced through attention-guided mechanisms and feature selection strategies are not merely aesthetic augmentations but reflect a deeper paradigm shift in the design philosophy of diagnostic AI systems. Modern clinical decision support tools are expected not only to be accurate but also to provide rationales for their predictions that are both intelligible and verifiable by clinical practitioners. The integration of interpretable outputs—whether through saliency maps, activation overlays, or attention heatmaps—aligns with this expectation and facilitates clinician trust, regulatory approval, and patient consent. These features also open the door for further clinical validation studies where the model's predictions can be compared directly with expert annotations in a qualitative framework, thereby supporting the validation of both sensitivity metrics and explanatory power. The model's demonstrated generalizability across different dataset splits and cross-validation folds underscores its robustness, yet it also raises questions about the portability of such architectures to entirely unseen datasets, including those with different imaging protocols, patient demographics, or lesion prevalence rates. Future investigations should explore domain adaptation and transfer learning techniques that can be harmonized with the current metaheuristic framework. This would allow the system to retain its adaptive capabilities while also

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

benefiting from pretrained knowledge acquired on external, large-scale, and perhaps even non-oral oncology datasets. In this context, metaheuristic control parameters could be dynamically tuned not only for architecture optimization but also for learning rate schedules, fine-tuning thresholds, and transfer layer selections. From a computational perspective, the resource-efficiency of the proposed system positions it favorably for real-world deployment scenarios. The architectural design choices, coupled with the optimization of model compression and pruning operations, result in an inference system that can function with minimal latency and modest hardware requirements. This is particularly relevant for deployment in community health clinics, mobile diagnostic units, and telemedicine platforms where access to high-performance computing resources may be constrained. However, such deployment scenarios introduce new variables, including variations in image capture devices, lighting conditions, and data resolution, all of which could influence model performance. Therefore, a future avenue of work could include the incorporation of data normalization pipelines and environment-aware calibration modules optimized via metaheuristics to enhance resilience under diverse acquisition conditions. Another noteworthy implication of the present study lies in its potential to inform longitudinal monitoring strategies. While the current model is primarily trained for diagnostic classification, its architecture could be readily adapted to support lesion tracking across time, enabling the monitoring of lesion progression, regression, or morphological change. This longitudinal perspective is critical for evaluating treatment response, recurrence probability, and progression from pre-malignant to malignant states. Integrating temporal data into the model's training pipeline, perhaps through recurrent modules or memory-based learning augmented by evolutionary scheduling, could extend the utility of the framework into the domain of prognostics. In the realm of data ethics and algorithmic equity, the model's design also addresses critical concerns about bias, transparency, and inclusivity. The optimization function was engineered to penalize outcome disparities across subgroups, thereby promoting a more equitable diagnostic experience for patients from diverse demographic backgrounds. This approach reflects a growing awareness that AI in healthcare must be fair by design, not as an afterthought. It also raises compelling questions about how optimization objectives are defined in future models—whether accuracy alone remains sufficient, or whether fairness, explainability, and robustness must be formalized as primary objectives in their own right. The clinical utility of the model is further reinforced by its capability for early-stage detection, an aspect that carries profound implications for patient prognosis and healthcare resource allocation. Detecting lesions at their earliest stages can significantly improve survival rates, reduce treatment costs, and minimize the physical and psychological burden on patients. The model's capacity to differentiate between high-risk pre-malignant lesions and benign conditions is particularly valuable in guiding clinical prioritization and resource triaging, especially in overburdened health systems. The operationalization of this functionality could be further enhanced by incorporating risk scoring algorithms and priority flags, allowing for real-time alerts in primary care settings. While the model performs strongly in offline environments, its translation to clinical practice will require rigorous validation through prospective trials, real-world testing, and user-centered design iterations. Questions of interface usability, integration with electronic health record systems, and the interoperability with laboratory information systems must be addressed. Moreover, collaboration with medical professionals will be essential to co-develop interpretability features, user feedback mechanisms, and decision thresholds that align with clinical intuition and workflows. These collaborative efforts could be facilitated through the implementation of continuous learning systems, where clinician feedback serves to update and refine the model over time, effectively creating a closed-loop diagnostic ecosystem. In terms of regulatory implications, the system's design aligns with many of the emerging frameworks for AI in medicine, including those emphasizing transparency, traceability, and robustness. However, more work is needed to define and test validation protocols that meet the stringent criteria of regulatory bodies across different jurisdictions. Real-world implementation must also consider cybersecurity, data governance, and auditability, all of which are essential to maintaining public trust and institutional accountability. The academic contributions of this work are not limited to the oral oncology domain. The architecture, optimization strategies, and design philosophies demonstrated here are transferable to other diagnostic domains, including dermatology, ophthalmology, gastroenterology, and pathology. The general principle of modular, metaheuristic-optimized, and interpretable diagnostic modeling can serve as a blueprint for developing AI systems that are both powerful and acceptable in clinical practice. Future research can explore the extension of this paradigm into multi-task settings, multi-modal input configurations, and hybrid human-in-the-loop systems. Furthermore, the methodology presented opens new avenues in optimization research itself. By embedding clinical constraints and ethical imperatives directly into the optimization function, the study demonstrates how metaheuristics can transcend mere parameter tuning and become instruments of value alignment. This suggests a broader research agenda in "ethical optimization," where objective functions reflect a multiplicity of goals including accuracy, equity, transparency, and usability. Looking ahead, the integration of real-time federated learning strategies with metaheuristic optimization could represent the next evolution of this work. Federated systems allow AI models to be trained across decentralized devices while

DOI: https://doi.org/10.54216/MOR.050203

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

preserving data privacy. When combined with metaheuristics that guide local learning rates, communication frequencies, and client sampling strategies, such systems could support continuously improving diagnostic models that learn from diverse clinical populations without centralizing sensitive patient data. This vision aligns with the future of decentralized, secure, and globally distributable healthcare AI. Finally, as artificial intelligence continues to permeate all aspects of healthcare delivery, the role of human-centered design in AI development becomes increasingly critical. Diagnostic models must not only perform well statistically, but must also be perceived as trustworthy, helpful, and empowering by their human users. The present study takes initial steps in this direction by embedding interpretability and equity into the optimization process. Future efforts can extend this approach by involving patients and providers in the model co-design process, developing explainability interfaces tailored to user literacy levels, and exploring emotional and psychological dimensions of AI-based diagnosis delivery. In conclusion, this expanded discussion reinforces the transformative potential of hybrid metaheuristic optimization in the development of diagnostic systems for oral cancer. The study provides strong evidence that optimization not only enhances technical performance but can also embed critical values such as fairness, transparency, and clinical relevance into the core of model design. Through its contributions to methodology, clinical practice, and future research, the present work establishes a framework that is both technically sound and ethically grounded—paving the way for next-generation AI systems that are intelligent, inclusive, and impactful.

4 Conclusion

The present study has explored the development and evaluation of a hybrid metaheuristic-optimized framework, The empirical outcomes of the study underscore the efficacy of the model in enhancing sensitivity, specificity, and generalization capability, while simultaneously offering interpretability through explainable intermediate representations. Moreover, the architecture was demonstrated to be scalable and lightweight, making it suitable for real-time deployment in both high-resource and resource-constrained settings. This opens pathways for broader clinical integration, particularly in mobile health applications, remote screening platforms, and point-of-care diagnostic systems. Beyond technical achievements, the study also emphasizes the ethical and societal implications of deploying intelligent diagnostic systems. Issues such as algorithmic fairness, data privacy, and clinical accountability were considered and addressed through careful model design and evaluation protocols. These considerations serve as essential pillars for responsible AI development in healthcare and form an integral part of the framework's long-term viability and acceptability in real-world settings. In conclusion, this study presents a significant advancement in the application of metaheuristic optimization to oral cancer diagnostics. The proposed system not only elevates the technical state of the art but also lays the groundwork for future research aimed at building intelligent, interpretable, and clinically deployable AI systems. As digital health technologies continue to evolve, such frameworks are poised to play a pivotal role in the early detection and management of cancer, ultimately contributing to improved patient outcomes and reduced disease burden on a global scaleThe broader implications of this study extend beyond the immediate technical achievements. By demonstrating that metaheuristic optimization can be effectively embedded into every stage of the diagnostic modeling pipeline—from data preprocessing to architecture selection and real-time inference—this work contributes to a growing consensus that bio-inspired algorithms are not merely auxiliary tools but core components of modern intelligent systems. Their adaptability, scalability, and cross-domain transferability position them as a foundational element in the development of next-generation healthcare technologies. Moreover, the emphasis placed on interpretability, clinical applicability, and computational efficiency ensures that the proposed framework aligns not only with academic objectives but also with the practical realities of healthcare systems. These systems often operate under constraints that include limited access to specialized diagnostic personnel, unequal distribution of technological infrastructure, and the need for cost-effective tools that can be deployed at scale. By prioritizing explainable outputs and real-time operability, the model responds directly to these challenges, offering a pathway for AI-assisted diagnostics to move from research settings into primary care, screening programs, and remote medicine environments frameworks like the one proposed here will become indispensable. They will inform not only diagnostics but also treatment planning, risk assessment, and patient monitoring. The future of healthcare will be increasingly defined by such systems—where computation, cognition, and clinical care intersect to improve outcomes, increase access, and empower both clinicians and patients alike.

DOI: https://doi.org/10.54216/MOR.050203

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

References

[1] Roozbeh Soltani and Amine Jaouadi. Optimization of Support Vector Machines by Meta-Heuristic Approach for Differential Diagnosis and Classification of Oral Diseases, August 2022.

[2] Akash Halder, Sugata Laha, Saptarshi Bandyopadhyay, Friedhelm Schwenker, and Ram Sarkar. A Metaheuristic Optimization Based Deep Feature Selection for Oral Cancer Classification. In Ching Yee Suen, Adam Krzyzak, Mirco Ravanelli, Edmondo Trentin, Cem Subakan, and Nicola Nobile, editors, *Artificial Neural Networks in Pattern Recognition*, volume 15154, pages 132–143. Springer Nature Switzerland, Cham, 2024. Series Title: Lecture Notes in Computer Science.

[3] T. Saraswathi and V. Murali Bhaskaran. Heuristic strategy using hybrid deep learning with transfer learning for oral cancer detection. *International Journal of Wavelets, Multiresolution and Information Processing*, 23(01):2450039, January 2025.

[4] Wenjing Wang, Yi Liu, and Jianan Wu. Early diagnosis of oral cancer using a hybrid arrangement of deep belief networkand combined group teaching algorithm. *Scientific Reports*, 13(1):22073, 2023. Publisher: Nature Publishing Group UK London.

[5] Qirui Huang, Huan Ding, and Navid Razmjooy. Oral cancer detection using convolutional neural network optimized by combined seagull optimization algorithm. *Biomedical Signal Processing and Control*, 87:105546, 2024. Publisher: Elsevier.

[6] Lei Zhang, Rongji Shi, and Naser Youssefi. Oral cancer diagnosis based on gated recurrent unit networks optimized by an improved version of Northern Goshawk optimization algorithm. *Heliyon*, 10(11), 2024. Publisher: Elsevier.

[7] Yirui Jiang, Trung Hieu Tran, and Leon Williams. Machine learning and mixed reality for smart aviation: Applications and challenges. *Journal of Air Transport Management*, 111:102437, August 2023.

[8] R. Rajkumar, Dınesh Valluru, Siva Satya Sreedhar P, N. Ramshankar, M. Sudha, and S. Navaneethan. Enhanced Jaya Optimization Algorithm with Deep Learning Assisted Oral Cancer Diagnosis on IoT Healthcare Systems. *Journal of Intelligent Systems & Internet of Things*, 11(2), 2024.

[9] Xiao Wei, Liu Chanjuan, Jiang Ke, Ye Linyun, Gao Jinxing, and Wang Quanbing. Convolutional neural network for oral cancer detection combined with improved tunicate swarm algorithm to detect oral cancer. *Scientific Reports*, 14(1):28675, 2024. Publisher: Nature Publishing Group UK London.

[10] Qirui Huang, Huan Ding, and Navid Razmjooy. Optimal deep learning neural network using ISSA for diagnosing the oral cancer. *Biomedical Signal Processing and Control*, 84:104749, 2023. Publisher: Elsevier.

[11] Dipali Wankhade, Chitra Dhawale, and Mrunal Meshram. Advanced deep learning algorithms in oral cancer detection: Techniques and applications. *Journal of Environmental Science and Health, Part C*, 43(2):133–158, April 2025.

[12] Saif Ur Rehman Khan and Sohaib Asif. Oral cancer detection using feature-level fusion and novel self-attention mechanisms. *Biomedical Signal Processing and Control*, 95:106437, 2024. Publisher: Elsevier.

[13] Narenthirakumar Appavu. Oral Cancer Histopathological Detection and Diagnosis Using Hybrid AI Deep Learning Methods. In *2025 3rd International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT)*, pages 2061–2067. IEEE, 2025.

[14] Ritu Gautam, Prableen Kaur, and Manik Sharma. Predictive Analytics for Early Cancer Detection Using Machine Learning and Generative AI. In *Generative Intelligence in Healthcare*, pages 150–175. CRC Press, 2025.

[15] Pullaiah Pinnika and K. Venkata Rao. Analysis of Oral Cancer Detection based Segmentation and Classification using Deep Learning Algorithms. In *International Conference on Computational Innovations and Emerging Trends (ICCIET-2024)*, pages 683–690. Atlantis Press, 2024.

DOI: https://doi.org/10.54216/MOR.050203
Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025

[16] K. J. Subha, M. Anto Bennet, Gaddam Pranay, Ketham Bharadwaj, and Polu Vikram Reddy. Analysis of deep learning based optimization techniques for oral cancer detection. In *2023 4th International Conference on Electronics and Sustainable Communication Systems (ICESC)*, pages 1550–1555. IEEE, 2023.

[17] El-Sayed M. El-Kenawy, Nima Khodadadi, Marwa M. Eid, Ehsaneh Khodadadi, Ehsan Khodadadi, Doaa Sami Khafaga, Amel Ali Alhussan, Abdelhameed Ibrahim, and Mohamed Saber. Improved cancer detection through feature selection using the binary Al Biruni Earth radius algorithm. *Scientific Reports*, 15(1):9483, 2025. Publisher: Nature Publishing Group UK London.

[18] Rinkal Shah and Jyoti Pareek. A Smartphone based Automated Primary Screening of Oral Cancer based on Deep Learning. *International Journal of Next-Generation Computing*, 15(3), 2024.

[19] Surbhi Vijh, Neetu Gupta, Sumit Kumar, and Priti Bansal. New bag of features using reinforcement Aquila optimization and weighted Bayesian Gaussian mixture modelling for dental images. *Expert Systems*, 41(6):e13453, June 2024.

[20] Lavanya Vemulapalli, Anantha Venkata Sai Kola, Chaitanya Chowdary Ravuri, and Akshara Kanagala. Strengthening oral cancer detection using WDCNN. In *2024 1st International Conference on Advances in Computing, Communication and Networking (ICAC2N)*, pages 1600–1605. IEEE, 2024.

[21] Masoumeh Motevalli, Madjid Khalilian, and Azam Bastanfard. Optimizing the Hybrid Feature Selection in the DNA Microarray for Cancer Diagnosis Using Fuzzy Entropy and the Giza Pyramid Construction Algorithm. *International Journal of Computational Intelligence and Applications*, 24(01):2450031, March 2025.

[22] Mohamed Ait Mehdi, Khadidja Belattar, and Feryel Souami. Recent metaheuristic algorithms for medical object localization using MSER detector in computer-aided diagnosis system. *Multimedia Tools and Applications*, 84(15):14433–14487, June 2024.

[23] Bandana Mahapatra and Anand Nayyar. Swarm intelligence and evolutionary algorithms for cancer diagnosis. In *Swarm intelligence and evolutionary algorithms in healthcare and drug development*, pages 19–44. Chapman and Hall/CRC, 2019.

[24] S. Pazhanirajan. Gazelle Optimized Visual Geometry Group Network with Resnet101 fostered Oral Squamous Cell Carcinoma Detection. *Journal of Intelligent Systems & Internet of Things*, 13(2), 2024.

[25] Mesfer Al Duhayyim, Areej Malibari, Sami Dhahbi, Mohamed K. Nour, Isra Al-Turaiki, Marwa Ismael Obayya, and Abdullah Mohamed. Sailfish Optimization with Deep Learning Based Oral Cancer Classification Model. *Comput. Syst. Sci. Eng.*, 45(1):753–767, 2023.

[26] Neel Bhattacharya, Ritwik Ghosh, and Soumya K Ghosh. Cnn-based classification for oral squamous cell carcinoma: An augmented dataset approach. *Biomedical Signal Processing and Control*, 70:102957, 2021.

[27] Shubham Garg, Mandeep Duhan, and Pooja Goel. Deep learning-based framework for detection of oral precancerous and cancerous lesions from photographic images. *Oral Oncology*, 126:105734, 2022.

[28] Abhishek Kumar, Rajat Sharma, and Ankit Chauhan. Metaheuristic-enabled deep feature selection and classification of oral cancer using cnn and svm. *Journal of King Saud University - Computer and Information Sciences*, 2023.

[29] Aman Singh, Meenakshi Rani, and Aashima Arora. A hybrid deep learning model for early detection of oral cancer using photographic images. *Health and Technology*, 10(6):1431–1441, 2020.

[30] Deepak Rao, Aditya Joshi, and Shruti Sharma. Diagnosis of oral cancer using transfer learning and convolutional neural networks. *Journal of Medical Systems*, 45:102, 2021.

[31] Ritesh Patel, Harsh Desai, and Het Thakkar. Lightweight deep learning model optimized with genetic algorithm for mobile-based oral cancer screening. *Computers in Biology and Medicine*, 144:105377, 2022.

DOI: https://doi.org/10.54216/MOR.050203

Received: June 06, 2025 Revised: August 23, 2025 Accepted: November 21, 2025