GA-Channel-Xception: Genetic Algorithm-Based Color Channel Selection for Deep Learning Detection of Oral Squamous Cell Carcinoma

João O. B. Diniz[1], Breno A. Tamanini[1], André F. Alvino[1]
Luana B. da Cruz[2], Josenildo C. da Silva[1], Omar A. C. Cortes[1],
Daniel L. Gomes Jr[1], Antônio O. de Carvalho Filho[3]

[1]Synthesis Lab – Fábrica de Inovação
Instituto Federal do Maranhão (IFMA)

[2]Laboratório de Inteligência Computacional Aplicada (LICA)
Universidade Federal do Cariri (UFCA)

[3]Programa de Pós-Graduação em Engenharia Elétrica
Universidade Federal do Piauí (UFPI)

joao.bandeira@ifma.edu.br

Abstract. *Oral Squamous Cell Carcinoma (OSCC) is a prevalent malignancy with high mortality due to late-stage diagnosis, driven by challenges in interpreting subtle histopathological image features. We introduce GA-Channel-Xception, a novel bioinformatics framework that integrates a Genetic Algorithm (GA) with the Xception convolutional neural network to enhance the detection of OSCC from histopathological images. Our approach leverages GA to optimize color channel selection across RGB, HSV, and LUV color spaces, identifying the most informative three-channel combination for classification. A tailored Xception model, trained on synthetic images derived from these channels, was evaluated on a publicly available OSCC histopathological dataset. The proposed method achieved an accuracy of 98.71% and an F1-score of 96.77%, outperforming conventional approaches. These results underscore the potential of integrating evolutionary algorithms with deep learning to improve automated histopathological analysis, offering a scalable tool for early OSCC diagnosis and precision oncology.*

1. Introduction

Oral cancer ranks among the most frequent malignant neoplasms worldwide, representing a serious public health concern [Murthy et al. 2025]. It is estimated that over 475,000 new cases are diagnosed annually on a global scale [WHO 2024]. In Brazil, oral cancer is among the most incident types in the 2023-2025 triennium [INCA 2023]. Among its variations, Oral Squamous Cell Carcinoma (OSCC) accounts for approximately 90% of cases, predominantly affecting men over 60 years of age with a history of smoking, alcohol consumption, and prolonged sun exposure [WHO 2024, Kumar and Nelson 2025].

Despite the importance of early diagnosis, most cases are identified at advanced stages, which reduces cure rates and significantly impacts patients' quality of

life [Murthy et al. 2025]. With a survival rate of around 50%, early detection is crucial for reducing mortality [Eckert et al. 2020]. Histopathological examination remains the gold standard for diagnosis, although it presents limitations such as inter- and intra-observer variability [Bisht et al. 2021].

In this scenario, advances in bioinformatics, particularly in digital pathology, have driven the use of computational methods for diagnostic assistance, with emphasis on deep learning models [Litjens et al. 2017]. These models have demonstrated high performance in the automatic classification of histopathological images, achieving results comparable to those of specialists [Carvalho et al. 2020, Ribeiro et al. 2024, Diniz et al. 2024a].

Visual image representation is a determining factor in classifier performance [Tamanini et al. 2025, Alvino et al. 2025]. Different color schemes can enhance or attenuate important morphological patterns. Thus, in addition to selecting the most suitable color space, choosing the most informative channels can significantly contribute to model effectiveness.

This work proposes a Genetic Algorithm (GA)-based approach to automatically select the optimal color channels from multiple schemes, thereby optimizing the input of histopathological images for OSCC detection using deep neural networks. The main contributions are:

- Proposal of a GA-based method for automatic channel selection from multiple color spaces;
- Integration of the selection with an optimized Xception architecture for OSCC classification.

2. Related Works

Several studies have investigated the application of deep neural networks for the automated detection of OSCC in histopathological images, aiming to reduce diagnostic errors and accelerate clinical screening processes. In [Deif et al. 2022], the authors propose a pipeline comprising feature extraction with different Convolutional Neural Networks (CNNs) such as VGG16, AlexNet, ResNet50, and InceptionV3, followed by feature selection using the Binary Particle Swarm Optimization (BPSO) algorithm, and classification with the XGBoost model. The best result was achieved by combining InceptionV3 with BPSO, reaching an accuracy of 96.3%.

In [Rahman et al. 2022], the approach is based on transfer learning, utilizing the AlexNet model to classify normal versus OSCC images. Results indicated an accuracy of 97.66% during training and 90.06% in testing. Meanwhile, the work of [Maia et al. 2024] evaluates different classification approaches, including traditional CNNs, transformers, and few-shot learning techniques. Among the tested models, DenseNet-121 achieved the best result with 91.91% accuracy.

The study by [Das et al. 2024] adopts an ensemble learning strategy, combining multiple CNNs (AlexNet, ResNet, Inception, and Xception) and selecting the top two to compose a final model. This approach achieved 97.88% accuracy. In [Raval et al. 2024], a modified ResNet is applied for classification, using data augmentation techniques and robust validation, reaching a validation accuracy of 96.72%.

The Histopath-DL-OC framework presented by [Murthy et al. 2025] conducts a systematic comparison among four convolutional architectures (VGG16, ResNet101, MobileNetV2, and InceptionV3). ResNet101 achieved the best performance with 97.21% accuracy.

Finally, the works of [Prado et al. 2025] and [Kumar and Nelson 2025] explore the same EfficientNet-B3 architecture for OSCC classification in histological images. Both studies reported 98.00% accuracy, reinforcing this architecture's high performance in biomedical contexts.

2.1. Related Works Evaluation

The literature highlights the predominant use of CNNs with transfer learning strategies and data augmentation, demonstrating robust accuracy and sensitivity results. However, we observe that none of these works explores the influence of color space or automatic channel selection as input variables.

Thus, existing works assume the RGB space as the optimal representation without evaluating alternative options. This limitation restricts potential image preprocessing optimization, a critical step for model construction.

Our work aims to address this gap by proposing a GA that automatically selects the three most informative channels from multiple color spaces. This selection is integrated into a deep learning pipeline with CNNs, empirically validating the impact of this choice on OSCC detection accuracy. The proposal expands discussions on automated preprocessing and input optimization for models used in digital pathology.

3. Proposed Materials and Methods

The method proposed in this work consists of four main stages. First, we define the database used for method construction. Next, we describe the application of the GA to select the most effective color channels for image representation. The third stage presents the classification process performed using an optimized Xception CNN. Finally, we describe the evaluation metrics used to validate and compare the robustness of the proposed method. Figure 1 illustrates these stages, which are detailed in the following sections.

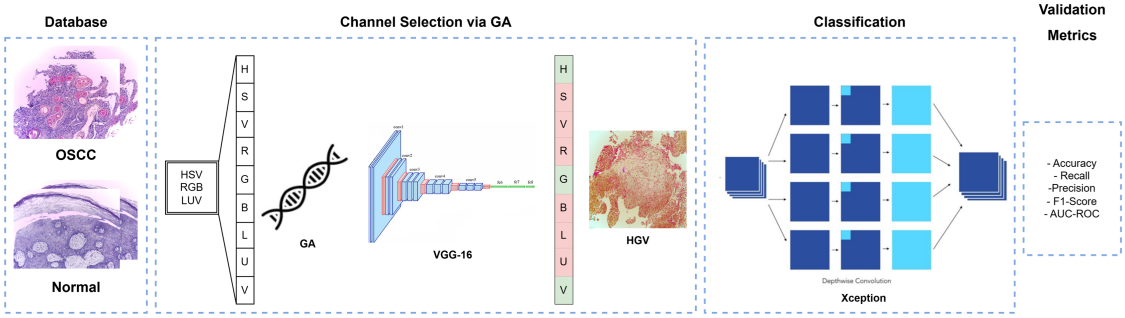

Figure 1. Graphical summary of the proposed method stages.

3.1. OSCC Database

This work uses the publicly available Histopathologic Oral Cancer Detection database from the Kaggle repository [Kebede 2022]. According to the authors, the dataset comprises 1,224 histopathological OSCC images stained with Hematoxylin and Eosin (H&E),

acquired by medical specialists from tissue slides of 230 distinct patients. Images were captured using a Leica ICC50 HD microscope with dimensions of 2040×1536 pixels. Representative samples of Normal and OSCC classes are shown in Figure 2.

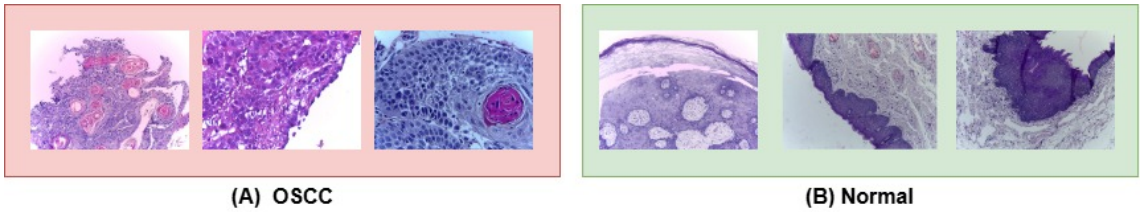

Figure 2. Database samples showing Normal and OSCC classes.

The dataset is pre-organized into three subsets with approximate proportions of 70% for training, 15% for validation, and 15% for testing. The training set includes data augmentation, totaling 2,435 Normal class images and 2,511 OSCC class images. The validation and test sets contain non-augmented images: 28 Normal and 92 OSCC for validation, and 31 Normal and 95 OSCC for testing.

We maintained this original division to preserve experimental integrity and enable fair comparisons with related work. A new split might artificially redistribute augmented images, potentially causing overlap of derived instances from the same original image and introducing evaluation bias.

Following database selection, we began the automatic channel selection phase to enhance image representation. Due to computational constraints in the selection and training processes, all images were resized to 224×224 pixels—a standard input resolution for CNNs, which, according to [Júnior et al. 2021, Diniz et al. 2024b, Diniz et al. 2024a], is suitable for deep learning pipelines.

3.2. Color Channel Selection via Genetic Algorithm

The representation of histopathological images plays a fundamental role in the performance of computer vision-based classification models. Studies such as [Alvino et al. 2025, Diniz et al. 2024a] demonstrate that the choice of color scheme can significantly impact classification results. Different color spaces may highlight distinct morphological and structural features in images, directly influencing the model's ability to discriminate between healthy tissues and cancerous lesions.

However, an approach combining channels from different color schemes could further enhance the extraction of intrinsic image features. Therefore, this work proposes an automatic mechanism for selecting the most informative color channels using a GA-based approach. The selection process considers nine channels from three distinct color spaces: RGB (Red, Green, Blue), HSV (Hue, Saturation, Value), and LUV (Luminance, U chromaticity, V chromaticity) [Gonzalez and Woods 2008].

3.2.1. Selected Color Schemes

The selection of these three color spaces (Section 3.2) is justified by their widespread use in biomedical image segmentation, classification, and contrast enhancement tasks.

RGB is commonly employed in most studies, as shown in Section 2. HSV and LUV were chosen for their superior ability to separate color and intensity components, which may facilitate the discrimination of structures and patterns in histopathological images. However, we emphasize that this approach is extensible, allowing for the inclusion of additional color schemes in future work. Due to the high computational cost associated with validating multiple color-space combinations, we limited our analysis to these three color spaces.

Thus, the original RGB image is converted to both HSV and LUV color spaces. From these spaces, we extract all nine individual channels. Each combination of three channels (forming a synthetic multichannel image) is evaluated to identify the configuration that maximizes the classification performance of our model.

3.2.2. Genetic Algorithm

The GA is a stochastic search technique inspired by principles of natural evolution, particularly genetic selection, crossover, and mutation [Shapiro 1999]. In this work, the GA serves as a combinatorial optimization mechanism to identify, among nine distinct color channels, the optimal combination of three channels that maximizes the classification performance of a CNN on a validation set. The objective function was the validation accuracy obtained after training a convolutional network (VGG16) using the selected channels as input.

Each population individual is represented by a three-element list corresponding to distinct channels chosen from the nine available options: `['rgb_r', 'rgb_g', 'rgb_b', 'hsv_h', 'hsv_s', 'hsv_v', 'luv_l', 'luv_u', 'luv_v']`. The initial population is randomly generated through sampling without replacement, ensuring each chromosome represents a unique and valid three-channel combination.

Individual fitness is evaluated through the performance of the VGG16 architecture with ImageNet pre-trained weights, adapted for binary classification, trained for 10 epochs on a reduced training set (50% of total data), using only the channels specified in the chromosome. The validation accuracy serves as the quality metric for each channel combination.

The main operators used to evolve the population across generations in which:

- **Crossover**: for each selected parent pair, the first two channels from the first parent are combined with the last channel from the second parent, avoiding duplicates. If needed, random channels are added until three distinct elements are obtained [Shapiro 1999].
- **Mutation**: with 30% probability, one of the three channels in the chromosome is replaced by a random (non-repeated) channel chosen from those not already present. It maintains genetic diversity and prevents premature population stagnation [Shapiro 1999].

For the GA implementation, we established several execution hyperparameters, including the number of generations, population size, and parent selection criteria. The specific hyperparameters were set as follows: Number of generations: 5; Population size:

6 individuals per generation; Parent selection: the two individuals with the highest accuracy in each generation are selected for reproduction. These values were chosen to obtain a trade-off between performance and computational feasibility.

As previously mentioned, to evaluate the performance of each channel combination, we employed a VGG-16 network on a reduced sample of the original dataset (50% of the training data). We validated it only on the validation set to prevent bias from the subsequent test set. These criteria were adopted due to the large number of possible combinations and the computational cost of executing the entire pipeline.

Upon completion of the GA execution, we obtained the optimal three-channel combination that best represents OSCC histopathological images. Subsequently, all dataset images were converted to this new scheme, generating synthetic images with the three selected channels.

3.3. Classification

At this stage, after converting all dataset images to the new channel scheme selected by the GA, they were used as input to a CNN based on the Xception architecture [Chollet 2017], adapted for OSCC histopathological image classification.

The choice of Xception is justified by its use of depthwise separable convolutions, which provide an efficient depth-to-computational-cost ratio, facilitating the extraction of subtle discriminative patterns even in datasets with sample limitations and class imbalance, as in our case study. Below, we detail the modifications made to Xception to improve its generalization capability:

- **Transfer learning**: Xception was initialized with ImageNet pre-trained weights, using its convolutional base as a feature extractor (frozen layers).
- **Additional dense layers**: at the output of the convolutional base, we added fully connected layers with 256 and 128 neurons, interspersed with `Dropout` layers of 0.5 and 0.3, respectively, for regularization and overfit reduction.
- **Weighted focal loss function**: we implemented a customized version of `Focal Loss` that, in addition to the focus parameter $\gamma = 2.0$ (default value used), also incorporates class weights (dynamically calculated based on the distribution of the dataset).
- **F1-score metric**: as the primary training evaluation metric, we used F1-score since the validation set has significant class imbalance and we seek the best trade-off between precision and recall.

With these modifications in place, Xception was trained and evaluated on the validation subset. Additionally, performance metrics were calculated on the test subset to validate the proposed method.

3.4. Validation Metrics

The proposed model was evaluated using several well-established classification metrics. Accuracy served as the general performance measure, indicating the proportion of correct predictions among all samples. Recall (or sensitivity) quantified the model's ability to correctly identify positive OSCC cases, while precision reflected the proportion of correct predictions among samples classified as positive. The F1-score provided the harmonic

mean between precision and recall. Furthermore, the Area Under the ROC Curve (AUC-ROC) measured the class separability. These metrics are standard for evaluating medical decision support systems [Powers 2020].

4. Results and Discussion

This section presents the experimental validation of the proposed method, including evaluations for each stage described in Section 3. We further discuss case studies, related work comparisons, and the main advantages and limitations of our approach.

All experiments were conducted on a computer with the following hardware and software specifications: Intel® Core™ i7 processor at 2.90 GHz, Nvidia RTX 4060 graphics card with 8 GB dedicated memory, 16 GB RAM, and Windows 11 Pro operating system.

The implementation was developed in Python using standard computer vision libraries including `NumPy`, `Pandas`, `OpenCV`, `Scikit-Learn`, and `Keras`. All code and experimental protocols will be made publicly available upon publication to ensure reproducibility.

4.1. Channel Selection Results

The GA execution yielded `['hsv_h', 'rgb_g', 'luv_v']` as the optimal channel combination, achieving 92.50% validation accuracy when evaluated with a VGG16 network trained on only 50% of the training subset. Figure 3 displays an example OSCC image in RGB space alongside the selected individual channels and the resulting synthetic 3-channel HGV composite image.

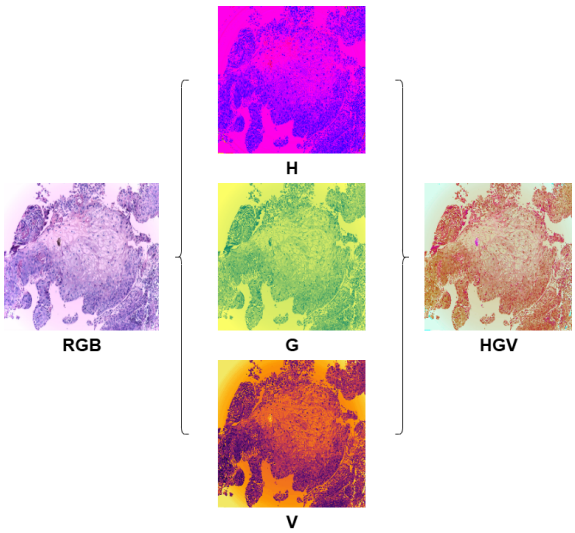

Figure 3. Results of automatic channel selection by the genetic algorithm.

These results demonstrate that combining information from different color spaces (HSV, RGB, and LUV) enhances discriminative representation of OSCC histopathological images, improving class separability between Normal and OSCC samples at the feature extraction stage.

Using this optimal combination, all dataset images were converted and used as input for training the Xception network. The network was trained as described in Section 3.3, using the Adam optimizer with a learning rate of 0.001, a batch size of 16, and 30 epochs.

4.2. Comparison Between Standard Color Spaces and Proposed Method

To validate the selection of cross-color-space channel combinations, we trained our complete proposed method (including the optimized Xception network with full training data) using standard RGB, HSV, and LUV color spaces, comparing their performance against our GA-selected channel combination. Table 1 presents these comparative results.

Table 1. Performance comparison between standard color spaces and the proposed GA-channel selection method.

Color Space	Accuracy	Recall	Precision	F1-Score	AUC-ROC
HSV	94.12%	91.75%	88.95%	90.33%	85.67%
RGB	93.25%	92.10%	89.56%	90.80%	87.45%
LUV	91.68%	89.23%	85.46%	87.25%	82.53%
GA-Channel-Xception	**98.71%**	**97.83%**	**95.74%**	**96.77%**	**94.88%**

The GA-selected channel combination coupled with the optimized Xception architecture demonstrated superior performance across all evaluation metrics. Achieving 98.71% accuracy, 97.83% recall, 95.74% precision, 96.77% F1-score, and 94.88% AUC-ROC, our proposed method significantly outperformed all standard color space configurations (RGB, HSV, and LUV). These results validate the effectiveness of the GA-based channel selection process, demonstrating that optimal channel selection directly impacts model performance for OSCC detection, establishing our approach as a robust and promising solution for histopathological image analysis.

4.3. Proposed Method vs Other CNN Architectures

We evaluated 11 established CNN architectures, including classical models (VGG16, ResNet, InceptionV3), optimized deep networks (DenseNet121), and modern efficient architectures (EfficientNet B0-B4). This selection represents a diverse range of architectural styles and complexity levels, allowing a comprehensive comparison with our Xception-based approach. All networks were trained using identical hyperparameters to Xception and the GA-selected channels. Table 2 presents the comparative results.

Table 2 demonstrates that all evaluated CNNs achieved satisfactory performance, with metrics exceeding 88% across all criteria. It confirms the effectiveness of our GA-based channel selection, which consistently enhanced each network's performance by identifying the most informative channel combinations for classification.

Among all architectures, our proposed GA-Channel-Xception method achieved superior results with 98.71% accuracy and 96.77% F1-score. These findings emphasize that while channel selection benefited all networks, the synergistic combination of optimal channels with the Xception architecture yielded the most significant improvements, establishing it as the optimal choice for our proposed framework.

Table 2. Performance comparison between different CNN architectures using GA-selected channels.

Architecture	Accuracy	Recall	Precision	F1-Score	AUC-ROC
VGG16	90.85%	88.34%	87.90%	88.12%	86.42%
DenseNet121	93.42%	91.15%	90.00%	90.57%	88.88%
InceptionV3	92.75%	90.92%	89.10%	89.99%	87.61%
ResNet50	94.10%	91.74%	91.10%	91.42%	89.94%
ResNet101	94.89%	92.55%	91.92%	92.23%	90.33%
EfficientNetB0	96.32%	94.66%	93.45%	94.05%	92.18%
EfficientNetB1	95.45%	93.42%	92.10%	92.75%	91.02%
EfficientNetB2	95.93%	94.00%	93.00%	93.49%	91.60%
EfficientNetB3	96.99%	95.66%	94.70%	95.18%	93.66%
EfficientNetB4	96.71%	95.02%	94.00%	94.50%	93.02%
GA-Channel-Xception	**98.71%**	**97.83%**	**95.74%**	**96.77%**	**94.88%**

4.4. Case Studies

As part of qualitative performance analysis, Figure 4 presents two representative examples from the test subset, both involving OSCC classification.

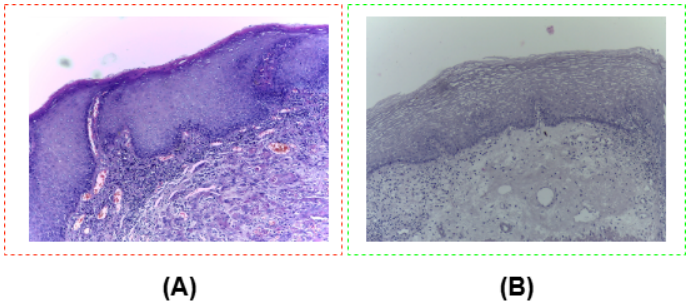

Figure 4. Case study analysis: (A) Normal sample misclassified as OSCC; (B) OSCC sample correctly identified.

Figure 4 (A) shows a Normal sample incorrectly classified as OSCC. This image displays high cellular density, nuclear overlap, and structural irregularities - features resembling malignant tissue that likely confused the model. Figure 4 (B) demonstrates a correctly identified OSCC case, exhibiting characteristic structural disorganization and cellular alterations typical of the disease. These examples highlight the model's challenges in distinguishing between healthy and malignant tissue when faced with morphological overlaps and structural transitions, illustrating the inherent complexity of histopathological diagnosis.

4.5. Comparison with State-of-the-Art

Table 3 provides a comprehensive comparison with leading literature on OSCC detection.

While multiple approaches achieved notable results (all exceeding 90% accuracy), as highlighted in Section 2, none of the existing works addressed image representation optimization. Our proposed method, which combines GA-based automatic channel selection

Table 3. Summary of related works with architectures, accuracy, recall and precision metrics.

Author(s)	Architecture(s)	Accuracy	Recall	Precision
[Deif et al. 2022]	Inception V3 + BPSO + XGBoost	96.30%	—	—
[Rahman et al. 2022]	AlexNet	90.06%	—	—
[Maia et al. 2024]	DenseNet-121	91.91%	91.93%	91.93%
[Das et al. 2024]	Ensemble CNN (ResNet + Xception)	97.88%	—	—
[Raval et al. 2024]	ResNet (Transfer Learning)	96.72%	—	—
[Murthy et al. 2025]	ResNet101	97.21%	—	—
[Prado et al. 2025]	EfficientNetB3	98.00%	—	—
[Kumar and Nelson 2025]	EfficientNetB3	98.00%	98.00%	98.00%
Proposed Method	**GA-Channel-Xception**	**98.71%**	**97.83%**	**95.74%**

with the Xception architecture, outperforms existing approaches with 98.71% accuracy, 97.83% recall, and 95.74% precision. These results demonstrate that intelligent channel selection has a direct and significant impact on classification performance.

Beyond confirming the effectiveness of established deep architectures, this study advances the state-of-the-art by proposing a novel approach that enhances CNN performance through the automated selection of color channels. Furthermore, it demonstrates strong potential for real-world clinical applications as a decision support tool for pathologists in digital histopathology.

5. Conclusions and Future Work

This study proposed a hybrid method for OSCC detection in histopathological images, combining automatic color channel selection via GA with Xception architecture classification. Our experimental results demonstrate the effectiveness of this approach, which outperforms existing methods in the literature, achieving 98.71% accuracy, 97.83% recall, and 95.74% precision. These findings highlight the crucial role of image representation in enhancing the performance of deep learning models in medical classification tasks, while also demonstrating the method's potential as a diagnostic support tool in pathology.

For future work, we propose: (1) optimization of Xception and GA hyperparameters, (2) exploration of model ensembles and Transformer-based architectures that have shown promise in computer vision tasks, and (3) comprehensive validation of model generalizability across additional public datasets to enhance clinical robustness and applicability. These directions would further strengthen the method's diagnostic reliability and practical utility in digital pathology workflows.

Acknowledgements

This work was carried out with the support of the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Code 001, Fundação de Amparo a Pesquisa do Maranhão (FAPEMA), and the Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq). Furthermore, we also acknowledge LLM use for spell-checking, grammar correction, and assistance in translating specific terms.

References

Alvino, A. F., Alves, E. S., Brito, S. S., Nascimento, V. T., da Cruz, L. B., Diniz, J. O., Souza Jr, L. O., da Silva, J. C., and Gomes Jr, D. L. (2025). Abordagem baseada em deep features para diagnóstico de câncer seroso de ovário em imagens histopatológicas. In *Simpósio Brasileiro de Computação Aplicada à Saúde (SBCAS)*, pages 401–412. SBC.

Bisht, S. R., Mishra, P., Yadav, D., Rawal, R., and Mercado-Shekhar, K. P. (2021). Current and emerging techniques for oral cancer screening and diagnosis: a review. *Progress in Biomedical Engineering*, 3(4):042003.

Carvalho, E. D., Antonio Filho, O., Silva, R. R., Araujo, F. H., Diniz, J. O., Silva, A. C., Paiva, A. C., and Gattass, M. (2020). Breast cancer diagnosis from histopathological images using textural features and cbir. *Artificial intelligence in medicine*, 105:101845.

Chollet, F. (2017). Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258.

Das, M., Dash, R., Mishra, S. K., and Dalai, A. K. (2024). An ensemble deep learning model for oral squamous cell carcinoma detection using histopathological image analysis. *IEEE Access*.

Deif, M. A., Attar, H., Amer, A., Elhaty, I. A., Khosravi, M. R., and Solyman, A. A. (2022). Diagnosis of oral squamous cell carcinoma using deep neural networks and binary particle swarm optimization on histopathological images: an aiomt approach. *Computational Intelligence and Neuroscience*, 2022(1):6364102.

Diniz, J. O., Ribeiro, N. P., Junior, D. A. D., da Cruz, L. B., de Carvalho Filho, A. O., Gomes Jr, D. L., Silva, A. C., and de Paiva, A. C. (2024a). Efficientxyz-deepfeatures: seleção de esquema de cor e arquitetura deep features na classificação de câncer de cólon em imagens histopatológicas. In *Simpósio Brasileiro de Computação Aplicada à Saúde (SBCAS)*, pages 82–93. SBC.

Diniz, J. O. B., Ribeiro, N. P., Dias Jr, D. A., da Cruz, L. B., da Silva, G. L., Gomes Jr, D. L., de Paiva, A. C., and Silva, A. C. (2024b). Anisotropicbreast-vit: Breast cancer classification in ultrasound images using anisotropic filtering and vision transformer. In *Brazilian Conference on Intelligent Systems*, pages 95–109. Springer.

Eckert, A. W., Kappler, M., Große, I., Wickenhauser, C., and Seliger, B. (2020). Current understanding of the hif-1-dependent metabolism in oral squamous cell carcinoma. *International journal of molecular sciences*, 21(17):6083.

Gonzalez, R. and Woods, R. (2008). *Digital image processing*. Pearson, Prentice Hall.

INCA (2023). Instituto nacional de câncer, estimativa 2023: Incidência de câncer no brasil. `https://www.inca.gov.br/publicacoes/livros/estimativa-2023-incidencia-de-cancer-no-brasil`. Accessed on: June 26. 2025.

Júnior, D. A. D., da Cruz, L. B., Diniz, J. O. B., da Silva, G. L. F., Junior, G. B., Silva, A. C., de Paiva, A. C., Nunes, R. A., and Gattass, M. (2021). Automatic method for classifying covid-19 patients based on chest x-ray images, using deep features and pso-optimized xgboost. *Expert Systems with Applications*, 183:115452.

Kebede, A. F. (2022). Histopathologic oral cancer detection using cnns. https://www.kaggle.com/datasets/ashenafifasilkebede/dataset. Accessed on: June 26. 2025.

Kumar, A. and Nelson, L. (2025). Enhancing oral squamous cell carcinoma detection using efficientnetb3 from histopathologic images. In *2025 International Conference on Multi-Agent Systems for Collaborative Intelligence (ICMSCI)*, pages 950–956. IEEE.

Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., Van Der Laak, J. A., Van Ginneken, B., and Sánchez, C. I. (2017). A survey on deep learning in medical image analysis. *Medical image analysis*, 42:60–88.

Maia, B. M. S., de Assis, M. C. F. R., de Lima, L. M., Rocha, M. B., Calente, H. G., Correa, M. L. A., Camisasca, D. R., and Krohling, R. A. (2024). Transformers, convolutional neural networks, and few-shot learning for classification of histopathological images of oral cancer. *Expert Systems with Applications*, 241:122418.

Murthy, A. S. R. C., Mercy, G., Prakash, L. J., and Bose, K. S. (2025). Histopath-dl-oc: Deep learning for oral cancer prediction from histopathology data. In *2025 International Conference on Inventive Computation Technologies (ICICT)*, pages 1190–1197. IEEE.

Powers, D. M. (2020). Evaluation: from precision, recall and f-measure to roc, informedness, markedness and correlation. *arXiv preprint arXiv:2010.16061*.

Prado, R. L. d., Marsicano, J. A., Frois, A. K., and Brancher, J. D. (2025). The use of machine learning to support the diagnosis of oral alterations. *Pesquisa Brasileira em Odontopediatria e Clínica Integrada*, 25:e240048.

Rahman, A.-u., Alqahtani, A., Aldhafferi, N., Nasir, M. U., Khan, M. F., Khan, M. A., and Mosavi, A. (2022). Histopathologic oral cancer prediction using oral squamous cell carcinoma biopsy empowered with transfer learning. *Sensors*, 22(10):3833.

Raval, D., Patel, A., Undavia, J. N., Shukla, A., and Patel, U. (2024). Oral cancer detection with convolutional neural networks and transfer learning: A resnet-based approach. In *International Conference on Data Analytics & Management*, pages 237–245. Springer.

Ribeiro, N. P., Teles, F. R., Diniz, J. O. B., da Cruz, L. B., Dias Jr, D. A., Braz Junior, G., de Almeida, J. D., and de Paiva, A. C. (2024). Improving colorectal cancer diagnosis using mirnet and inceptionv3 on histopathological images. In *Brazilian Conference on Intelligent Systems*, pages 321–334. Springer.

Shapiro, J. (1999). Genetic algorithms in machine learning. In *Advanced course on artificial intelligence*, pages 146–168. Springer.

Tamanini, B. A., Sousa, V. G., Rodrigues, L. P., Oliveira, D. M., Dias, C. X., da Cruz, L. B., Diniz, J. O., and Júnior, L. O. S. (2025). Classificaçao de carcinoma endometrioide de ovário por transformaçao de esquema de cor e radiomics em imagens histopatológicas. In *Simpósio Brasileiro de Computação Aplicada à Saúde (SBCAS)*, pages 68–79. SBC.

WHO (2024). World health organization: Oral cancer. https://www.who.int/news-room/fact-sheets/detail/oral-health. Accessed on: June 26. 2025.