

Deep Learning-Based Detection and Classification of Dental Conditions

Mrs. Priyata Mishra, Jaikumar Dewangan, Avijit Agrawal, Bhupendra Dewangan

Dept. of Computer Science and Engineering SSIPMT, Raipur, C.G., India

DOI: <https://doi.org/10.51583/IJLTEMAS.2026.150100078>

Received: 28 January 2026; Accepted: 02 February 2026; Published: 10 February 2026

ABSTRACT

Dental conditions, particularly calculus and caries, are among the primary causes of poor oral health. They frequently result in pain, discomfort, infections, and occasionally even tooth loss. To prevent these problems from getting worse, early detection is vital. In many places, though, access to dental specialists is still restricted. In order to automatically identify and categorize four crucial aspects of oral health—caries, calculus, discoloration, and healthy teeth—this study proposes a deep learning approach. By using readily available intraoral images, the model removes the need for radiographic methods up to an extent, allowing for affordable and non-invasive screening. The system uses Transfer Learning with a ResNet-based Convolutional Neural Network (CNN), which provides excellent feature extraction and quick learning even with smaller datasets. We trained the model extensively with pre-processed intraoral images and fine-tuned the higher convolutional layers. In tests on the validation dataset, the model reached a classification accuracy of 97.35 percent. This shows it can accurately tell the difference between healthy and diseased dental states with great precision and low variance 0.34 percent across the cross-validation tests. We can expand the community's access to early detection through integrating them into applications. This framework marks a major breakthrough in the use of deep learning models in preventive dentistry with its rapid, scalable, and accurate screening that improves patient care and public health outcomes.

Keywords; Dental Conditions, Deep Learning, CNN, AI.

INTRODUCTION

Regardless of age or location, dental caries, plaque, and calculus affect a significant portion of the population. This makes oral diseases some of the most common illnesses [1]. Despite being mostly preventable, these conditions often remain untreated due to a lack of early detection and insufficient access to dental care facilities, particularly in developing and rural areas [2]. Poor oral health can lead to pain and tooth decay, and it's also connected to other diseases like diabetes and heart disease. This highlights the importance of early detection [3]. Effective preventative care also depends on consistently monitoring healthy teeth. Most dental exams rely on clinical inspections and if having more symptoms goes to radiographic imaging, both require trained specialists, equipment, and in-person visits. This reliance creates significant barriers to quick diagnoses, as many people delay care until they feel discomfort, leading to costly and invasive treatments later. In addition, regional and socioeconomic differences in dental resources increase the gap in access to oral healthcare [4]. We need to create affordable, accessible diagnostic tools that can be used outside of traditional clinical settings. Recent advancements in Artificial Intelligence (AI), mainly Deep Learning (DL), have greatly improved medical image analysis. Among Deep Learning methods, Convolutional Neural Networks (CNNs) stand out for their exceptional ability to automatically extract spatial and visual features from complex images [5]. Architectures such as ResNet allow for deeper networks to be trained without losing performance, increasing the accuracy of visual interpretation without performance degradation [6]. Additionally, Transfer Learning, where a model trained on huge dataset is adjusted for a specific area, can significantly improve performance, even with limited medical data, speeding up development [7]. This research uses CNN-based Transfer Learning to identify and classify four major dental categories—caries, calculus, discoloration, and healthy teeth—from non-radiographic, intraoral images. The focus on real-image data, which can be captured by simple cameras or smartphones,

enhances the system's applicability for community use and improves accessibility for remote dental assessments. Including the 'healthy teeth' category is important for building a foundation for preventive oral health monitoring. With a classification accuracy of 97.35%, the proposed model shows great promise for real-world use as a digital diagnostic tool for broad preventive oral screening.

LITERATURE REVIEW

The field of dental diagnostics has witnessed significant evolution, driven by the need for AI-based image analysis systems capable of augmenting expert-level visual interpretation. Historically, dentists have relied on manual inspection and radiographic imaging for assessing caries and other oral abnormalities. Although clinically effective, these methods suffer from inherent subjectivity and limited scalability for mass screening purposes [8].

Evolution of Computer Vision in Dental Diagnostics

Several initial studies explored applying conventional computer vision methods in dentistry. Earlier, the image processing techniques focused on thresholding and edge detection for caries detection; however, these approaches lacked robustness against variations in illumination, texture, and angle typical of real-world intraoral images [9]. The subsequent introduction of classical machine learning algorithms, particularly Random Forest and Support Vector Machines (SVMs) classifiers, improved the detection of dental calculus and enamel anomalies but still necessitated labour-intensive, manual feature extraction, limiting their potential for wide-scale, automated systems [10].

The Deep Learning Paradigm and Architectural Advance

Deep Learning (DL) marked a fundamental paradigm shift. CNNs have demonstrated the capability to achieve human-level accuracy in classifying numerous medical conditions from images—including skin lesions, fundus images for diabetic retinopathy, and dental pathologies visible on radiographs [11]. Crucially, CNNs can automatically learn intricate, multi-level features directly from raw image pixels without the need for domain-specific feature engineering, a vital factor for managing the high variability of intraoral images captured under non-clinical conditions. He et al. (2015) introduced the Residual Network (ResNet), a deep CNN architecture that addresses the vanishing gradient problem through the use of skip connections, making it possible to train networks with hundreds of layers effectively [6]. This innovation is credited for the architecture itself [12], which has yielded breakthrough performance in complex image classification tasks by enabling the network to learn residual mapping.

Table 1. Dataset Distribution

Category	Number of Images	Percentage
Caries	2,381	33.92%
Calculus	1,307	18.62%
Healthy Teeth	1,501	21.38%
Discoloration	1,831	26.08%
Total	7,020	100%

Transfer Learning for Domain Specialization

Transfer Learning has since extended these architectural benefits into healthcare, particularly where high-quality, domain-specific datasets are limited. By initializing models with ImageNet-trained weights and then fine-tuning on dental datasets, the model inherits powerful low-level visual features (e.g., edges, corners) learned from

millions of general images. This process significantly accelerates convergence and enhances To ensure consistent input quality and increase model robustness, a structured preprocessing pipeline was developed. In order to maintain equal input size, each of the image dimensions was resized to 224x224. To normalize the distribution of the data to simplify further processing and accelerate convergence, we normalized pixel intensity by normalizing it using a division by 255 in $[0, 1]$.

To address class imbalance and enhance model generalization, a selective class-adaptive data augmentation strategy was used. To create the illusion of mirrored orientations, all images were subjected to horizontal flipping. Strong augmentation (e.g., $\pm 15\%$ – 40% brightness, 0.7 – $1.4X$ contrast adjustment) was used for the minority class (Calculus). The majority of classes received light augmentation with a $\pm 15\%$ brightness variation limit. This adaptive approach ensured a more balanced representation and reduced overfitting.




Fig. 1.Flowchart of Proposed Methodology

Contemporary AI Applications in Dental Screening

Recent dental research directly supports the application of advanced CNNs and Transfer Learning using non-radiographic imagery. Studies leveraging ResNet50 and similar architectures have reported high accuracy in detecting caries and calculus using intraoral photographs [14]. A few pioneering works have attempted to classify healthy versus diseased teeth, showing strong potential for self-assessment applications via smartphone cameras [15]. This trend indicates a growing recognition that surface-level pathologies can be reliably identified without X-rays, paving the way for ubiquitous screening tools. Building upon this foundation, the present study integrates these architectural and methodological advances into a robust, multi-class system specifically targeting the foundational conditions of caries, calculus, discoloration, and the crucial 'healthy teeth' status from real intraoral images. This approach simultaneously addresses the clinical need for high diagnostic accuracy and the public health imperative for improved accessibility.

METHODOLOGY

The methodology details the refined, deep learning-based approach specifically designed for the automated and accurate classification of the four core dental health states using solely intraoral images. This system is engineered to provide an accessible diagnostic alternative to traditional clinical inspections.

Data Acquisition and Preprocessing

The study made use of a combined dataset of 7,020 intraoral dental images in JPG format that were obtained from Mendeley Data and Kaggle, two publicly accessible and peer-reviewed repositories [16] [17]. All photos were meticulously and manually annotated by skilled dental professionals to guarantee high-quality labels.

To ensure consistent input quality and increase model robustness, a structured preprocessing pipeline was developed. In order to maintain equal input size, each of the image dimensions was resized to 224x224. To normalize the distribution of the data to simplify further processing and accelerate convergence, we normalized pixel intensity by normalizing it using a division by 255 in [0, 1]. To address class imbalance and enhance model generalization, a selective class-adaptive data augmentation strategy was used. To create the illusion of mirrored orientations, all images were subjected to horizontal flipping. Strong augmentation (e.g., $\pm 15\%$ – 40% brightness, 0.7 – $1.4\times$ contrast adjustment) was used for the minority class (Calculus). The majority of classes received light augmentation with a $\pm 15\%$ brightness variation limit. This adaptive approach ensured a more balanced representation and reduced overfitting.

Model Architecture and Training Strategy

The proposed system uses a transfer learning technique with ResNet50 architecture [7]. The architecture is composed of two primary parts: a feature extraction backbone, and a custom classification head.

Table 2. Model Summary

Layer (type)	Output Shape	Param #	Connected to
input_layer_9 (InputLayer)	(None, 224, 224, 3)	0	-
get_item_12 (GetItem)	(None, 224, 224)	0	input_layer_9[0][0]
get_item_13 (GetItem)	(None, 224, 224)	0	input_layer_9[0][0]
get_item_14 (GetItem)	(None, 224, 224)	0	input_layer_9[0][0]
stack_4 (Stack)	(None, 224, 224, 3)	0	get_item_12[0][0], get_item_13[0][0], get_item_14[0][0]
add_4 (Add)	(None, 224, 224, 3)	0	stack_4[0][0]
resnet50 (Functional)	(None, 7, 7, 2048)	23,587,712	add_4[0][0]

global_average_pooling2d (GlobalAveragePooling2D)	(None, 2048)	0	resnet50[0][0]
dense_8 (Dense)	(None, 256)	524,544	global_average_pooling2d[0][0]
dropout_4 (Dropout)	(None, 256)	0	dense_8[0][0]
dense_9 (Dense)	(None, 4)	1,028	dropout_4[0][0]

Model Initialization and Custom Head

The backbone was initialized using ResNet50 weights pretrained on ImageNet and set as a non-trainable feature extractor in the first stage. A custom classification head was added comprising a Global Average Pooling 2D layer for spatial dimensions reduction, a ReLU activation function, L2 regularization ($\lambda = 1e-4$), a Fully Connected (Dense) layer with 256 units, a Dropout layer (rate = 0.5) for regularization, and an Output layer with 4 neurons and Softmax activation for multi-class classification. In order to address class imbalance, the Scikit-learn “balanced” class weighting approach was used during training.

Training Methodology: A two phase transfer learning method was used to enable good adaptability of features.

Phase 1 — Feature Extraction: All the convolutional layers of ResNet50 were pre-trained, while only the custom head was trained over 25 epochs by freezing ResNet Base as feature extractor with Adam optimizer (learning rate = $1e-4$) and Sparse Categorical Cross-Entropy loss. This step allowed the classifier to learn how to map pre-trained generic features into dental disease classes.

Phase 2 — Fine-Tuning: The ResNet50 backbone's initial 100 layers were set to a frozen state. A very low learning rate (LR= $1e-5$) was used to fine-tune the entire model for 12 additional epochs. This controlled process avoids catastrophic forgetting and allows the network to adapt its high-level features to the specific visual characteristics of the intraoral images. Early stopping and learning rate reduction callbacks were utilized to ensure stable convergence and prevent overfitting.

Evaluation Metrics and Validation

A 3-fold cross-validation scheme was used to make sure each image was used as a test sample once in order to guarantee a reliable and objective assessment. The training subset was further split into an 80/20 split for training and validation, respectively, within each fold. Accuracy, Balanced Accuracy, Macro F1-Score, Precision, and Recall were among the metrics used to evaluate the model's performance, in addition to a thorough Confusion Matrix analysis. In order to ensure generalization across data splits, the reported results are the average performance across all three folds.

RESULTS AND DISCUSSION

The rigorous evaluation of the proposed ResNet50-based deep learning model, employing the two-phase Transfer Learning strategy, yielded exceptional performance metrics, strongly validating the system's potential for automated, non-radiographic dental screening.

Performance Evaluation and Model Consistency

The performance of the model was assessed through a 3-fold cross-validation approach on the 7,020 images dataset. The summarized performance of all folds is shown in Table II. The system exhibited strong stability across partitions, with very small variation ($\sigma \leq 0.34\%$), suggesting generalization performance more robust

rather than over-fitted to a single data split. An average overall Accuracy of 97.35% verifies the model as a very accurate diagnostic method.

Table 3. Performance metrics

Metric	Fold 1	Fold 2	Fold 3	Mean $\pm$ STD
Test Accuracy	97.48%	97.69%	96.88%	97.35% $\pm$ 0.34%
Balanced Accuracy	97.65%	97.85%	97.03%	97.51% $\pm$ 0.35%
F1-Score (Macro)	97.72%	97.90%	97.16%	97.59% $\pm$ 0.32%

Per-Class Performance and Clinical Reliability

Detailed per-class metrics (Figure 3) confirm the model's clinical utility. The model achieved near-perfect performance for key categories: Calculus detection reached a Recall of 0.9985 (no false negatives), and the classification of Healthy Teeth attained an F1-Score of 0.9957 (Precision 0.9960, Recall 0.9953). The high precision and recall for caries (0.9461 and 0.9887 respectively) underscore the model's reliability in identifying the most prevalent pathology.

Table 4. Per-class Performance Metrics

Class	Precision	Recall	F1-Score
Caries	0.9461	0.9887	0.9669
Calculus	0.9805	0.9985	0.9894
Healthy	0.9960	0.9953	0.9957
Discoloration	0.9882	0.9181	0.9519
Accuracy	—	—	0.9735
Macro Avg	0.9777	0.9751	0.9760
Weighted Avg	0.9742	0.9735	0.9733

Comparative Context and Architectural Impact

The achieved accuracy of 97.35% represents a substantial advancement over recent literature in dental AI. Our framework demonstrates a notable 5.5-7.6% improvement in accuracy for four-class classification. This superior performance is directly attributable to the combined benefits of the ResNet50 architecture and the optimized two-phase Transfer Learning strategy.

Error Analysis and Implications

Analysis of the aggregated Confusion Matrix (Figure. 3 - *Normalized Confusion Matrix*) revealed a critical

finding: approximately 79.6% of all misclassifications occurred between the Discoloration and Calculus classes. This 2.65% inter-class confusion rate is clinically understandable, as early calculus deposits and certain pigmentation patterns share visual similarities under varied illumination. This error is manageable given the high overall accuracy, but it provides a clear direction for future refinement, such as utilizing more advanced colour normalization techniques or domain-specific loss functions.




Fig. 4. Confusion Matrix

CONCLUSION

In this paper, we have demonstrated high performance and feasibility of a deep-learning based approach to automatic classification of four basic oral-health states (caries, calculus, discoloration, healthy teeth) employing only readily obtainable intraoral images. The model obtained a remarkable average classification accuracy of 97.35% on the test set, when by using Transfer Learning based ResNet architecture. While the high diagnostic precision and use of prevalent RGB images emphasize how AI can support decision making in clinician settings, chances to expand dependable dental screening into non-clinician environments are feasible.

The correct distinction of the 'healthy teeth' state is especially important, since it forms a basis for efficient preventive treatments. This mechanism has low implementation cost and is generally available, supporting the proactive management of oral health. Next, through collaboration with mobile platforms, large-scale community-based dental wellness screening programs will be facilitated and early disease diagnose and oral care education can benefit people from all over the world.

REFERENCES

1. A. Qayyum *et al.*, "Dental caries detection using a semi-supervised learning approach," *Scientific Reports*, vol. 13, no. 1, 0 2023, doi: 10.1038/s41598-023-27808-9. [Online]. Available: <https://pubmed.ncbi.nlm.nih.gov/36639724/>
2. M. E. Northridge, A. Kumar, and R. Kaur, "Disparities in Access to Oral Health Care.," *Annual Review of Public Health*, vol. 41, no. 1, pp. 513–535, 0 2020, doi: 10.1146/annurev-publhealth-040119-094318. [Online]. Available: <https://pubmed.ncbi.nlm.nih.gov/31900100/>
3. M. S. Lipsky, T. Singh, G. Zakeri, and M. Hung, "Oral Health and Older Adults: A Narrative Review.," *Dentistry Journal*, vol. 12, no. 2, p. 30, 0 2024, doi: 10.3390/dj12020030. [Online]. Available: <https://www.mdpi.com/2304-6767/12/2/30>

4. K. Griffith, J. Bor, and L. Evans, "The Affordable Care Act Reduced Socioeconomic Disparities in Health Care Access.," *Health Affairs*, vol. 36, no. 8, pp. 1503–1510, 0 2017, doi: 10.1377/hlthaff.2017.0083. [Online]. Available: https://www.researchgate.net/publication/318733499_The_Affordable_Care_Act_Reduced_Socioeconomic_Disparities_In_Health_Care_Access
5. Y. Xu, R. Quan, W. Xu, Y. Huang, X. Chen, and F. Liu, "Advances in Medical Image Segmentation: A Comprehensive Review of Traditional, Deep Learning and Hybrid Approaches.," *Bioengineering (Basel, Switzerland)*, vol. 11, no. 10, p. 1034, 0 2024, doi: 10.3390/bioengineering11101034. [Online]. Available: <https://www.mdpi.com/2306-5354/11/10/1034>
6. K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016. [Online]. Available: https://openaccess.thecvf.com/content_cvpr_2016/html/He_Deep_Residual_Learning_CVPR_2016_paper.html
7. M. Raghu et al., "Transfusion: Understanding Transfer Learning for Medical Imaging," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 32, 2019. [Online]. Available: https://papers.nips.cc/paper_files/paper/2019/hash/eb1e78328c46506b46a4ac4a1e378b91-Abstract.html
8. Schwendicke, F et al. "Artificial Intelligence for Caries Detection: Value of Data and Information." *Journal of dental research* vol. 101,11 (2022): 1350-1356. doi:10.1177/00220345221113756. [Online]. Available: <https://pubmed.ncbi.nlm.nih.gov/35996332/>
9. A. M. Shervedani, H. Khodadadi, and S. I. Mousavian, "Development a computer-aided diagnosis system for dental caries detection applying radiographic images," *Computers in Biology and Medicine*, vol. 196, Part C, p. 110966, 2025, doi: 10.1016/j.combiomed.2025.110966. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0010482525013186>
10. Choi J. Comparative effectiveness research in observational studies. *J Periodontal Implant Sci*. 2018 Dec;48(6):335-336. <https://doi.org/10.5051/jpis.2018.48.6.335>. [Online]. Available: <https://jpis.org/DOIx.php?id=10.5051/jpis.2018.48.6.335>
11. A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, no. 7639, pp. 115–118, 2017. [Online]. Available: <https://www.nature.com/articles/nature21056>
12. Y. Chen et al., "Deep Transfer Learning for Multi-Pathology Classification in Dental Radiographs," *IEEE Access*, vol. 7, pp. 107797–107806, 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8713964>
13. Mohammad-Rahimi, Hossein et al. "Deep learning for caries detection: A systematic review." *Journal of dentistry* vol. 122 (2022): 104115. doi:10.1016/j.jdent.2022.104115 [Online]. Available: <https://pubmed.ncbi.nlm.nih.gov/35367318/>
14. Kang S, Shon B, Park EY, Jeong S, Kim EK. Diagnostic accuracy of dental caries detection using ensemble techniques in deep learning with intraoral camera images. *PLoS One*. 2024 Sep 6;19(9):e0310004. doi: 10.1371/journal.pone.0310004. PMID: 39241044; PMCID: PMC11379315. [Online]. Available: <https://pmc.ncbi.nlm.nih.gov/articles/PMC11379315/>
15. Takahama, Ricardo Kenji et al. "Accuracy of smartphone photographs for detecting active carious lesions in orthodontic patients." *Brazilian oral research* vol. 39 e069. 8 Sep. 2025, doi:10.1590/1807-3107bor-2025.vol39.069 [Online]. Available: <https://pmc.ncbi.nlm.nih.gov/articles/PMC12419182/>
16. S. Salman, "Oral Diseases Dataset," Kaggle, 2023. Accessed: Jan. 2025. Available: [Online]. Available: <https://www.kaggle.com/datasets/salmansajid05/oral-diseases>
17. "Dental Health Image Collection," *Mendeley Data*, 2023. [Online]. Available: <https://data.mendeley.com/datasets/>