

Transfer Learning and Data Augmentation to Improve the Efficacy of a Periodontal Disease Screening Model

Atharv Gupta¹, Aryan Anumula²

Received October 30, 2024

Accepted September 11, 2026

Electronic access November 15, 2026

Oral Diseases pose a major health burden to the world. Diagnosis of Oral Diseases and in particular Periodontal Disease or Gingivitis is expensive, and uneven and inaccessible to a lot of people. Deep learning approaches can help in quick screening and initial diagnosis of Periodontal disease that can help accessibility and lower costs. Previous studies on utilizing deep learning for periodontal disease have focussed on high quality medical images taken in clinical settings which may not be possible in rural settings. The goal of this work is to design an AI-powered periodontal health analyzer that makes dental health analysis and preliminary diagnosis more accessible utilizing non-clinical images. To achieve this, a two-stage deep learning pipeline was developed: a mouth validation model to detect and filter valid oral cavity images, followed by a periodontal disease classifier to screen for signs of gingivitis. Although a custom convolutional neural network performed well on the training data, it overfitted on the test data. This limitation was resolved by applying targeted data augmentation techniques to expand training diversity, while optimal performance was achieved through transfer learning by fine-tuning a pre-trained VGG-16 model. Ultimately, the mouth validation model achieved an accuracy of 96.1% on the validation set, and the optimized periodontal disease classifier scored a validation accuracy of 85.9% alongside a test set accuracy of 84.0%. Overall, the work demonstrated that data augmentation and transfer learning can be leveraged to build a high quality periodontal disease classifier using Oral images.

Keywords: Periodontal Disease Screening, Gingivitis Detection, Deep Learning, Convolutional Neural Networks (CNN), Transfer Learning, Data Augmentation, VGG-16.

Introduction

Oral diseases are a major health burden affecting around 3.5 billion people worldwide according to the World Health Organization. The population affected by Oral diseases continues to rise globally due to urbanization, fluoride deficiency, excessive intake of food with high sugar content and poor access to oral health care¹. There is a critical need for scalable diagnostic frameworks². Dental Health is vital to a person's overall health. Often dental diseases lead to more serious diseases of the respiratory system, digestive system and the cardiovascular system. Emerging clinical research has increasingly highlighted the bidirectional link between chronic oral infections and systemic inflammatory disorders^{3,4}.

Treatment and diagnosis for oral health conditions is expensive and usually not part of universal health coverage (UHC). Consequently, financial barriers remain a primary obstacle to early dental intervention on a global scale^{5,6}. A policy brief by UCLA Center for Health Policy Research by Nadaret Pourat et al shows that there is shortage of dentists in some parts of California and a lot of state dentists don't provide care for low income residents who rely on MediCAL⁷. Access to adequate

oral healthcare is uneven throughout the world. It is estimated that 2.4 billion people have either little or no access to oral health care⁸. According to the American Dental Association, around 23% of U.S. adults don't have dental benefits⁹. Even if you do have insurance, it will only cover part of the cost of more expensive treatments and it might not cover some procedures at all. So there is a big need for low cost and efficient diagnosis of dental diseases which is accessible to everyone.

Gingivitis is a common and mild form of gum disease, and is also called periodontal disease. It's one of the earliest stages of gum disease. Oral cleanings can help successfully manage gingivitis. Gingivitis, if left untreated can lead to a more severe form called periodontitis. Over time, chronic periodontitis results in irreversible bone destruction and finally tooth loss⁴. A dentist can help diagnose gingivitis by looking for explicit symptoms like bleeding, gum recession etc or by looking at dental x-rays. Gingivitis symptoms aren't always obvious in the early stages and so you may have it and not know it¹⁰.

Problem Statement and Rationale

Machine Learning and Deep Learning are subsets of Artificial intelligence. ML and DL can help predict various outcomes by training on provided datasets using a variety of algorithms, such as linear regression, logistic regression, and neural net-

¹ Cupertino High School

² Evergreen High School

works. Neural networks try to mimic the human brain and have interconnected artificial neurons that receive and analyze data signals. Artificial intelligence (AI) has been useful in oral disease diagnosis, predicting prognosis, and developing patient-specific treatment strategies¹¹⁻¹³.

ML can help in quick diagnosis of periodontal disease which can increase accessibility and lower costs. ML based solutions can help individuals as primary screening tools and prompt them to seek dental care and they can also help the clinicians read and detect disease from images more quickly, freeing up more appointment time and hence lower costs and increase accessibility.

There has been wide research on using machine learning and deep learning models for dental disease detection utilizing clinically available images and features. These include microscopy images, endoscopic images, radiographs or X-ray images or professional grade intra-oral images. Patil et al explored the use of artificial intelligence in the diagnosis of Oral Diseases and did a literature survey of applications of AI in various dental conditions¹⁴. Ozden et al. examined the use of 3 machine learning algorithms (SVMs, Decision tree, and a neural network) to identify and classify periodontal disease. Clinical data was obtained from a total of 150 patients. The systems classified the data into six types of periodontal conditions¹⁵. From dental plaque microscopy images, Aberin and Goma used the convolutional neural networks as the classifier and utilized the AlexNet architecture to classify the images as healthy or unhealthy. The overall accuracy of the system was found to be 75.5%¹⁶. Balaei et al. demonstrated the automatic detection of periodontitis using intra-oral images¹⁷. Kim et al. reported that using CNNs may lessen the obligations of dental radiologists who interpret images by successfully evaluating periodontal bone loss from panoramic radiographs using a CNN system¹⁸. Chau et al trained a deep learning model on 567 high quality intra oral photographs using SLR labelled per site along gingival margins¹⁹. Further primary research by Krois et al., Lee et al., and Casalegno et al. demonstrated the efficacy of deep learning on pristine clinical radiographs and intraoral imagery²⁰⁻²². Additional recent studies have continuously validated AI performance in clinical settings²³.

A comprehensive review of the existing research reveals a major gap in how artificial intelligence has been deployed for periodontal evaluation. The research is heavily focussed on using high-cost clinical imaging methods like radiographic modalities, panoramic scans, and professional-grade intraoral lenses^{14,15,18,20}. These methods have strong reliance on specialized clinical infrastructure. Models trained exclusively on clinical datasets suffer from severe performance degradation when they are used on non-clinical smartphone images due to real world noise from variable illumination, camera angles, and lower resolution of smartphone imaging^{19,24}. This highlights the need for robust and accessible periodontal screen-

ing systems that are built to handle the real world noise of consumer images while maintaining high accuracy and performance. This will be particularly useful for dental disease screening in rural and inaccessible regions. More recent literature has started focusing on non-clinical and accessible consumer images. Mania L et al built models and achieved sensitivity of 77.6% and specificity of 93.8% on oral health screening of children using smartphone photos. However this study was focussed on children aged 7-12 and gingivitis progression may be different in children²⁴. Li, W., Liang, Y. et al. used Multi-Task learning to automate screening of gingivitis and dental caries based on oral RGB photos and outperformed general-purpose CNNs²⁵. They used Multi-task learning for the two tasks of gingivitis classification and localization and obtained 87% AUC on gingivitis classification. While this study attempted mobile or RGB photo screening, they lacked robust handling of fully unconstrained and noisy consumer imagery. Thanh et al. focussed on using deep learning on smartphone photos for dental caries²⁶.

Our research aims to leverage low-cost, unconstrained non-clinical RGB photos to build a highly accurate periodontal disease screening system which can do automated feature extraction and can compensate for real world noise coming from variable lighting, camera angles and lower resolution.

Significance, Objectives, and Scope

The literature survey above highlights how deep convolutional networks and attention mechanisms are leveraged extensively in vision-based medical diagnostics. However, applying these techniques and models directly to low-resolution oral consumer images often results in low performance which can be attributed to training overfitting and out of domain data. This area of periodontal disease detection is relatively less explored. This study addresses this limitation by investigating how common architectures like VGG-16, when combined with data augmentation, can be successfully adapted to combat this limitation. Our work is focussed on taking non-clinical oral images which are more varied and lower quality than medical imaging for periodontal disease classification. The objective of our work is to create an oral health analyzer that can help screen for oral diseases from images given by the user. It is intended to be a screener and analyzer tool which can help folks detect early oral diseases. In this work, we focussed on periodontal disease but the work can be extended to other oral health conditions.

Methodology Overview

The paper is organized as follows. The Methodology section outlines the proposed methodology, including problem formulation, the dataset used and the modeling strategies which in-

clude details on data augmentation techniques, and transfer learning. The Results section presents the experimental results for each model iteration including the final proposed architecture. The Discussion section is a discussion on model behavior, analysis of techniques used and potential limitations of this study. The Conclusion section finally presents the Conclusion providing a summary and outlines directions for future research.

Methodology

Problem Formulation

The overall task of periodontal screening is broken down into a two-stage binary classification system. Both tasks take non-clinical digital images as input. The first classifier (Mouth Validation Classifier) takes an arbitrary RGB image of size $180 \times 180 \times 3$ and classifies it into a binary output: mouth or non-mouth. The second classifier (Periodontal Disease Classifier) takes a validated mouth image and classifies it into either presence of gingivitis or healthy. The core problem solved by this system is to filter out non-oral images while detecting signs of early periodontal disease in oral images.

The proposed methodology is outlined in Figure 1.

Input Dataset

The Mouth Validation model had a total of 4201 images that were positive (images of mouth) and 9876 images that were negative (images of non-mouths). The Periodontal Disease Classifier model had a total of 2034 images that were positive (images of gingivitis) and 1987 images that were negative. Figure 2 shows some samples of images with gingivitis and without gingivitis.

The input dataset was prepared in the following way.

1. **Data Collection & Sources:** The dataset is a collection of non-clinical RGB oral images which was obtained from crowdsourcing platforms and open-source dental repositories. The images in this collection were all captured using various consumer smartphone cameras. These images therefore have real-world non-clinical peculiarities including uneven illumination, variable focus, and multiple angles.
2. **Labeling Criteria:** Ground-truth labels for periodontal disease were present in the input sources and were determined by clinicians based on explicit clinical indicators such as local gum redness, swelling, stippling loss and bleeding along the margins. To ensure high quality, we chose only the input data that is already cross-verified against standard diagnostic guidelines.
3. **Potential Biases and Limitations:** The dataset suffers from various selection biases. The images collected don't




Figure. 1 Proposed Methodology

have a balanced representation across age groups and geographic locations, and skin colors. Inflammation in gingivitis manifests differently and in varying degrees depending on oral mucosal pigmentation. As a result, models trained on lighter tissue may not be accurate on darker tissues. Future improvements to this work would require a more diverse dataset to ensure good diagnostic performance across all demographics and age groups.

4. **Dataset Splitting:** The dataset was split into a training set (70%) used for training, a validation set (15%) used for hyperparameter tuning and checkpoint selection, and a held-out test set (15%) used for final evaluation.

Modeling Approaches

Mouth Validation

The Mouth Checker, our Image Screening Model, acts as the first line of defense in our Oral Health Analyzer. It filters uploaded images, ensuring only oral images proceed for further analysis. With precision levels exceeding 95%, this model effectively identifies and discards non-dental images, enhancing



Figure. 2 Sample dataset images illustrating the diversity of input data, including positive mouth captures with gingivitis, healthy mouth tissues, and negative non-mouth images filtered out by the initial validation model.

the overall accuracy. It is an off the shelf CNN model that aims to improve the screening process, facilitating efficient and reliable oral health assessments. Specifically, Keras Sequential models were trained using the Tensorflow python library.

Periodontal Disease Classifier

The Periodontal Disease Classifier is the core component of our Oral Health Analyzer, designed to detect a range of targeted dental diseases. This model scans oral cavity images for signs of gingivitis. We started with a CNN model and then used various techniques to improve the accuracy of this model.

CNN Model

We trained a CNN model for the Periodontal Disease classifier. We experimented with various CNN model configurations. Each configuration had interspersed convolution layers followed by max pooling layers and then a final feed forward layer. We had to hit the right balance of layers and parameters. Increasing the number of layers and parameters in the model to very high led to overfitting of the model on training data and while decreasing it to a small number of parameters led to low performance.

The CNN architecture progressively extracts hierarchical features from the input images. The initial convolution layers capture local features like edges and color gradients for gum lines. Pooling layers help to downsample the dimensions ensuring small shifts in mouth placement do not affect the classification results. Later convolutional blocks combine the local extracted features to understand complex semantic patterns which represent disease features like gum tissue inflammation or texture changes. Finally, these are all flattened and concatenated as input to the feed forward layers for final output.

The proposed configuration included 3 convolution layers followed by max pooling and a final fully connected network

with a hidden layer of 128 neurons followed by the output layer. The model accepts input images of size $180 \times 180 \times 3$. It first applies rescaling to normalize input pixel values into a floating point number in the range $[0, 1]$. It then feeds the input into three 2D convolutional blocks.

- The first Conv2D layer uses 16 convolution filters and preserves the input dimensions (180×180). It is followed by a MaxPooling2D that downsamples the image to 90×90 .
- The second Conv2D layer uses 32 convolution filters and outputs a $90 \times 90 \times 32$ feature map. This is followed by a MaxPooling2D layer that reduces it to $45 \times 45 \times 32$.
- The third Conv2D layer uses 64 convolution filters and outputs a $45 \times 45 \times 64$ feature map. Another MaxPooling2D layer downscales it to $22 \times 22 \times 64$.
- The resulting feature maps are flattened and concatenated into a 30,976-dimensional vector.
- This is the input to a fully connected layer with 128 hidden units followed by an output layer for binary classification). Overall, the model contains about 3,988,898 trainable parameters.

We train the model for 10 epochs. We experimented with various learning rates, batch sizes, and dropout rates. The candidate model used a learning rate of 0.0001 with Adam optimizer, batch size of 32 and also a dropout rate of 0.2 for regularization.

Data Augmentation

Data Augmentation is a technique used in machine learning to improve the accuracy and generalizability of ML models by creating new training data from existing samples. It helps decrease overfitting by introducing diversity into the training set. Augmented data is derived from original data with minor transformations. However, synthetic data is typically generated using architectures like Generative Adversarial Networks (GANs). Data augmentation was applied exclusively to the training dataset to prevent data leakage and ensure an unbiased and real evaluation on the test set.

For the periodontal disease classifier, non-clinical oral RGB images captured via smartphone cameras have several real world variations. These variations like awkward shooting angles, shifting head orientations, uneven ambient lighting, and camera motion blur need to be smoothen out. To make the model generalize well across these real world variations, we applied three types of data augmentation techniques.

1. **Geometric Transformations like Rotation, Flipping, Stretching, Zooming, and Cropping:** Users capture

photos of their mouths from various angles and distances. A standard CNN is sensitive to global and local positions and this transformation helps it learn from symptoms appearing in different parts of the mouth. Random rotation, zoom, and cropping forces the model to learn location-insensitive features and identify inflammation regardless of its global location or size within the image. Horizontal flipping simulates pictures of teeth and gum structures taken from different sides of the oral cavity.

2. **Color Transformations like Hue Shifts, Brightness, and Contrast Adjustments:** Non-clinical smartphone images have drastic variations in lighting depending on indoor lighting, camera flash type, and specifications of cameras used. Changes in brightness or color balance may result in the model misclassifying healthy tissue as periodontal (or vice versa). Randomly modifying the hue, brightness, and contrast helps to balance the color distribution in the training data. This helps the model to recognize true periodontal redness and swelling whether the image was taken at night or under bright daylight.
3. **Kernel Transformations like Blurring and Sharpness adjustments:** Non-clinical smartphone images often get out of focus or blurred due to hand movement during capture. Applying data augmentation that randomly blurs or sharpens the training images simulates the real world. As a result, the model does not rely on image sharpness, and it can identify inflammation when analyzing blurry images.

These augmentation techniques help in creating a more balanced distribution of data from the initial noisy and real world smartphone images. Figure 3 shows some samples of images created using data augmentation:



Figure. 3 Data Augmented Images

Transfer Learning

Transfer Learning is a machine learning technique in which a model trained for 1 task or dataset is used to improve model performance on a different task which is related. As described in section B on Input Dataset, we had a limited amount of labeled data even after data augmentation. The primary insight of transfer learning in this context is that a general purpose model pretrained on a large dataset of images can help improve the accuracy of the specific task of periodontal image classification task.

We considered VGG16, ResNet50 and InceptionNetV2 as candidate pre-trained models to leverage transfer learning. VGG16 is a CNN-based deep model and uses 3×3 convolutions. Resnet uses residual or skip connections to help train deeper networks. InceptionNetV2 is an efficient architecture that captures multi-scale features with reduced computation. We chose to leverage transfer learning from the VGG16 model²⁷ to improve our candidate model performance primarily due to its simpler structure and ease of debugging and integration. VGG16 model's uniform and linear stack of layers is much easier to understand and debug. It is much easier to apply transfer learning by freezing the base layers of VGG16 because of clear separation between the convolutional and the feed forward layers in VGG16. Highly optimized models like Inception and ResNet use advanced techniques (e.g., special convolutions for dimensionality reduction, complex branching) to achieve computational efficiency. However, these models may make the feature maps less direct for the periodontal disease task where precise spatial localization of subtle damage is important.

VGG16 is a CNN architecture and model that specializes in image classification and image recognition for a wide variety of images²⁷. VGG16 can be applied to determine whether an image contains certain items, animals, plants and more. VGG16 is a large network and has about 138 million parameters.

It is composed of 16 layers of neurons which work to process image information into a latent representation which is then used for image classification tasks. VGG-16 uses convolution layers with a kernel size of 3×3 and a stride of 1. It has maxpool layers of 2×2 . The number of channels is doubled after every 2 or 3 Convolution layers to identify multiple feature maps. Each feature map represents in vector space useful features learnt from the input image to help with the classification task. The various feature maps are then concatenated with each other in the flattening layer which is a representation of all the features learnt from the image in latent space. This is then fed to three fully connected dense layers, followed by a softmax operation for output. The softmax function outputs the probability distribution over all the output classes (1000).

VGG16 network is pre-trained on over 1 million images in the ImageNet visual database, and is able to classify images




Figure. 4 VGG model architecture

into 1000 categories. The model achieves accuracy of 92.7 percent in the top-5 test.




Figure. 5 Final Model architecture

We modified the VGG16 model for our Periodontal classification task by replacing the output layer with just a single output logit and using the standard binary cross entropy function instead of softmax. We had to apply rescaling of the input images from $180 \times 180 \times 3$ to $224 \times 224 \times 3$ as the VGG16 model accepts images of that size. We fine-tuned the pretrained weights of the VGG16 model with data augmented training data. Upon analysis, it became apparent that incor-

poration of a pretrained CNN model (VGG16) using transfer learning led to an improvement in the model's accuracy.

Performance Evaluation Metrics

We used a confusion matrix to evaluate the binary classification models for the two tasks. The four cells in the matrix represent the total number of True Positive (TP), True Negative (TN), False Positive (FP), and False Negative (FN) examples.

We derived four metrics from the confusion matrix.

1. Accuracy: $(TP+TN) / (TP+TN+FP+FN)$
2. Precision: $TP / (TP+FP)$
3. Recall (Sensitivity): $TP / (TP+FN)$
4. F1-score: $2 \times (Precision+Recall) / Precision \times Recall$

Results



Figure. 6 Training and Validation Accuracy for original CNN (b) Training and Validation Accuracy with data augmentation (c) Training and Validation Accuracy with data augmentation

The original CNN model achieved a training set accuracy of 85% after 10 epochs but validation set accuracy of only around 76%. This indicated overfitting and we experimented with regularization schemes to deal with this. However the validation set accuracy didn't improve (figure 6a). We then started training the models with data augmentation. The performance of the validation set and the training set were similar on the augmented datasets with accuracy still approaching close to 76% after 14 epochs of training (figure 6b). So, the overfitting issue observed with the original CNN model was resolved. We then employed transfer learning using pre-trained VGG models modified for this task and used the data augmented training data and observed a significant improvement in accuracy on the validation set of 85.9% as shown in Figure 6c. We then used a held out test (15% of the total dataset, $n=603$ images) to measure the generalization capability of the final optimized model. The model achieved an overall test accuracy of 84.08%.

The confusion matrix for the test set is shown in the table below.

Metric / Category	Predicted Positive (Gingivitis)	Predicted Negative (Healthy)	Total Actual
Actual Positive (Gingivitis)	256 (TP)	46 (FN)	302
Actual Negative (Healthy)	50 (FP)	251 (TN)	301
Total Predicted	306	297	603 (Test Set)

The model predicted 256 true positives and 251 true negatives, with 50 false positives and 46 false negatives resulting in Precision of 83.66%, Recall of 84.77%, and an F1-score of 84.21%. The test set results confirmed the model's generalization to an unseen held out set. It demonstrated good performance in screening non-clinical real world oral images for periodontal disease.

To validate whether the performance improvement from transfer learning was statistically significant, a paired t-test was conducted. The results confirmed that a CNN model trained with data augmentation and leveraging transfer learning from a pre-trained VGG-16 model produced statistically significant increase in accuracy ($p < 0.01$) compared to just training a similar CNN from scratch.

Discussion

The experimental results highlight the efficacy of combining data augmentation and transfer learning for non-clinical oral disease screening.

In this section, we analyze the experimental results to understand the trade-offs between the baseline architecture and techniques like data augmentation and transfer learning when dealing with non-clinical RGB images.

- **Baseline CNN and Overfitting:** The original CNN model achieved a high training accuracy (85%) but stagnated on the validation set (76%). This was due to the high variance in non-clinical images from inconsistent lighting and background noise and the model being trained from scratch on a relatively small dataset. Hence, the model memorized specific training set noise patterns instead of learning generalized features of periodontal disease.
- **Impact of Data Augmentation:** Adding the data augmented images including the geometric, color, and kernel transformations helped to resolve the overfitting issue since the training data variance due to non-clinical features was smoothened out and the model could generalize on the validation set. However, data augmenta-

tion did not significantly increase the absolute validation accuracy past 76%. The reason is that augmentation increases diversity of background image features, however, it cannot invent missing semantic features that are needed to distinguish periodontal disease from similar looking healthy oral images. .

- **Transfer Learning:** Fine-tuning a pretrained VGG model led to very significant performance improvement of reaching an 85.9% validation accuracy. VGG-16 has been pre-trained on over a million ImageNet images and its convolutional layers have learnt discriminative features for edge, texture, and color detection. So when we do transfer learning, these pre-existing layers in the initial part of the network allow the model to effectively detect periodontal disease from noisy oral cavity images.

Although the overall accuracy of 84% on the test set is good, the model has both False positives and False Negatives. It sometimes misclassifies benign gum irritation or natural pigmentation as periodontal disease. Sometimes, it misses to detect early-stage or periodontal images under poor lighting. So this can at best be used as a screening tool and we cannot rely solely on AI models built on top of smartphone camera images as a replacement for professional clinical diagnosis.

Conclusion

In this study, we developed and evaluated an AI-powered, two-stage oral health screening system designed to make preliminary periodontal disease screening accessible using consumer non-clinical oral images. The Mouth validation model successfully filtered oral images with an accuracy of 96.1%. For the periodontal screening model, the original CNN model initially suffered from overfitting. This was rectified by integrating synthetic and methodical data augmentation. Leveraging transfer learning by fine-tuning a pre-trained VGG-16 architecture improved overall performance with validation set accuracy of 85.9% and a held out test set accuracy of 84.0%.

This study demonstrates that deep learning can be effectively used for dental health screening and analysis purposes using non-clinical consumer images. Existing literature on dental image analysis have focussed on clinical grade images and this study helps in bridging the gap between high-end clinical datasets and noisy real-world consumer images. Transfer learning and data augmentation are effective techniques that help in generalizing and augmenting the performance on non-clinical noisy images. We validated our hypotheses that data augmentation helps in resolving overfitting and transfer learning boosts accuracy across real world image sets.

Future research of this work should focus on increasing overall volume and diversity of images across various demographic groups and skin-tones. The dataset was collected from

the internet and suffers from potential selection and demographic biases. This may influence real world performance as the model may not be equally sensitive across diverse skin tones. Furthermore, the model performance can be potentially improved from implementing multi-task learning to simultaneously classify and segment the periodontal disease markers in images. We can also deploy the model in an easy to access cross-platform mobile application.

This work offers a low-cost easily accessible screening mechanism for dental health analysis. This cannot replace a professional clinical diagnosis, but is a good resource for underserved communities with limited accessibility to dental clinics to seek timely clinical intervention. Early dental screening through an accessible deep learning model is one of the several ways in which we can reduce global healthcare disparities.

References

- 1 J. Wang, "Oral Health," World Health Organization, 2023. [Online]. Available: <https://www.who.int/health-topics/oral-health>. [Accessed: Aug. 2024].
- 2 P. E. Petersen and H. Ogawa, "The global burden of periodontal disease: towards integration with chronic disease prevention and control," *Periodontology* 2000, vol. 60, no. 1, pp. 15-39, 2012.
- 3 M. Sanz et al., "Periodontitis and cardiovascular diseases: Consensus report," *Journal of Clinical Periodontology*, vol. 47, no. 3, pp. 268-288, 2020.
- 4 D. F. Kinane, P. G. Stathopoulou, and P. N. Papapanou, "Periodontal diseases," *Nature Reviews Disease Primers*, vol. 3, no. 1, pp. 1-14, 2017.
- 5 M. R. Mathur et al., "Universal Health Coverage: A Unique Policy Opportunity for Oral Health", *Journal of Dental Research*, 94(3 Suppl): 3S-5S, 2015.
- 6 M. A. Peres et al., "Oral diseases: a global public health challenge," *The Lancet*, vol. 394, no. 10194, pp. 249-260, 2019.
- 7 N. Pourat et al., "Report identifies barriers to accessing dental care for low-income Californians," UCLA Center for Health Policy Research, 2021. [Online]. Available: <https://www.uclahealth.org/news/release/report-identifies-barriers-accessing-dental-care-low-income>. [Accessed: Aug. 2024].
- 8 FDI World Dental Federation, "Smile Around the World," 2023. [Online]. Available: <https://www.fdiworlddental.org/smile-around-world>. [Accessed: Aug. 2024].
- 9 Health Policy Institute, "National Trends in Dental Care Use, Dental Insurance Coverage, and Cost Barriers," American Dental Association, 2024. [Online]. Available: https://www.ada.org/-/media/project/ada-organization/ada/ada-org/files/resources/research/hpi/national_trends_dental_use_benefits_barriers_2024.pdf. [Accessed: Aug. 2024].
- 10 Cleveland Clinic, "Gingivitis," 2023. [Online]. Available: <https://my.clevelandclinic.org/health/diseases/21482-gum-periodontal-disease>. [Accessed: Aug. 2026].
- 11 T. Shan, F. R. Tay, and L. Gu, "Application of Artificial Intelligence in Dentistry," *Journal of Dental Research*, vol. 100, no. 3, pp. 232-244, 2021.
- 12 A. Esteva et al., "A guide to deep learning in healthcare," *Nature Medicine*, vol. 25, no. 1, pp. 24-29, 2019.
- 13 E. J. Topol, "High-performance medicine: the convergence of human and artificial intelligence," *Nature Medicine*, vol. 25, no. 1, pp. 44-56, 2019.
- 14 S. Patil et al., "Artificial Intelligence in the Diagnosis of Oral Diseases: Applications and Pitfalls," *Diagnostics*, vol. 12, no. 5, p. 1029, 2022.
- 15 F. O. Ozden, O. Özgönenel, B. Özden, and A. Aydogdu, "Diagnosis of Periodontal Diseases Using Different Classification Algorithms: A Preliminary Study," *Nigerian Journal of Clinical Practice*, vol. 18, 2015.
- 16 S. T. A. Aberin and J. C. de Goma, "Detecting periodontal disease using convolutional neural networks," in *IEEE 10th Int. Conf. Humanoid, Nanotechnology, Information Technology, Communication and Control, Environment and Management (HNICEM)*, Baguio City, Philippines, 2018, pp. 1-6.
- 17 A. T. Balaei, P. de Chazal, J. Eberhard, H. Domnisch, A. Spahr, and K. Ruiz, "Automatic detection of periodontitis using intra-oral images," in *39th Annu. Int. Conf. IEEE Engineering in Medicine and Biology Society (EMBC)*, Jeju, Korea, 2017, pp. 3906-3909.
- 18 T. S. Kim, C. Obst, S. Zehaczek, and C. Geenen, "Detection of bone loss with different X-ray techniques in periodontal patients," *Journal of Periodontology*, vol. 79, no. 7, pp. 1141-1149, 2008.
- 19 R. C. W. Chau et al., "Accuracy of Artificial Intelligence-Based Photographic Detection of Gingivitis," *International Dental Journal*, vol. 73, no. 5, pp. 724-730, 2023.
- 20 J. Krois et al., "Deep learning for the radiographic detection of periodontal bone loss," *Scientific Reports*, vol. 9, no. 1, p. 8495, 2019.
- 21 J. H. Lee et al., "Detection and diagnosis of dental caries using a deep learning-based convolutional neural network algorithm," *Journal of Dentistry*, vol. 77, pp. 106-111, 2018.
- 22 F. Casalegno et al., "Caries Detection with Near-Infrared Transillumination using deep learning," *Journal of Dental Research*, 98(11), 1227-1233, 2019.
- 23 Mutawa et al., "Deep Learning Applications for Dental-Disease Classification Using Intraoral Photographic Images: Current Status and Future Perspectives," *AI*, 7(3), 85, 2026.
- 24 L. Mania, K. Nanobashvili, T. Manjavidze, M. Benashvili, and I. Ashtamadze, "Diagnostic accuracy of intraoral mobile photography for oral health screening in children: a pilot study," *BMC Oral Health*, vol. 25, no. 1, p. 1144, 2025.
- 25 W. Li, Y. Liang, X. Zhang, et al., "A deep learning approach to automatic gingivitis screening based on classification and localization in RGB photos," *Scientific Reports*, vol. 11, p. 16831, 2021.
- 26 Mai Thi Giang Thanh et al., "Deep Learning Application in Dental Caries Detection Using Intraoral Photos Taken by Smartphones", *Applied Sciences*, vol. 12, no. 11, p. 5504, 2022.
- 27 K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in *International Conference on Learning Representations (ICLR)*, 2015.