Deep Learning for Dental Problems and Their Detection: DentalAI

Aabhesh Singh Shekhawat, Harshit Singh Bisht, Anjali Singh

Abstract

Oral diseases such as dental caries, calculus, gingivitis, and mouth ulcers are among the most commonly occurring health conditions worldwide and often remain undiagnosed in their early stages due to lack of awareness and limited access to dental care. Early detection of these diseases is essential to prevent serious complications and improve overall oral health. In this study, an automated dental disease detection system named DentalAI is proposed using deep learning for the classification of common oral diseases from dental images. The system is developed using the ResNet34 convolutional neural network architecture with the help of transfer learning to improve learning efficiency and accuracy.

A publicly available oral disease image dataset was used for model training and evaluation. The dataset was carefully preprocessed and enhanced using image augmentation techniques such as rotation, flipping, zooming, and brightness adjustment to improve generalization and reduce overfitting. The model was trained to classify four disease categories: calculus, dental caries, gingivitis, and mouth ulcers. Experimental evaluation was carried out on a test set consisting of 620 dental images.

The proposed model achieved an overall classification accuracy of 80.48%. Class-wise performance analysis showed strong detection capability for gingivitis and mouth ulcers, while moderate confusion was observed between clinically similar conditions such as calculus and gingivitis. The confusion matrix and performance metrics confirmed that the system learned meaningful disease-specific visual features and demonstrated stable predictive behavior across all classes.

The results indicate that the proposed system can serve as an effective AI-assisted preliminary dental screening tool, especially in rural and underserved areas where access to dental specialists is limited. Although the system is not intended to replace professional dental diagnosis, it provides a fast, low-cost, and accessible solution for early-stage oral disease detection. With further improvements in dataset size, class balance, and model optimization, the performance of the system can be enhanced further. This work highlights the strong potential of deep learning techniques in the development of smart and accessible dental healthcare systems.

Keywords:

DentalAI , Oral Health, Deep Learning, ResNet34, Dental Disease Detection, Dental Caries, Calculus, Gingivitis, Mouth Ulcers, Medical Image Analysis, Convolutional Neural Networks (CNN), Artificial Intelligence in Healthcare, Early Diagnosis, Automated Dental Screening, Smart Healthcare System, Digital Dental Consultation, Preventive Oral Care.

1.Introduction

Oral health plays a vital role in a person's overall well-being, confidence, and quality of life. A healthy mouth not only allows individuals to eat, speak, and smile comfortably but also protects the body from various infections and systemic diseases. However, despite it's importance, dental health is often neglected by a large section of the population, especially in developing regions. Many people avoid visiting dentists due to fear, lack of awareness, high treatment costs, busy lifestyles, or poor access to healthcare facilities. As a result, small dental problems gradually turn into serious oral diseases that cause pain, discomfort, infection, and even permanent loss of teeth.

Among the most common and widely occurring oral conditions are calculus, dental caries, gingivitis, and mouth ulcers. Dental caries (tooth decay) is one of the most widespread oral diseases globally and develops due to the destruction of tooth enamel by harmful bacteria. Calculus, also known as tartar, forms when plaque hardens on the teeth and leads to severe gum problems if not removed in time. Gingivitis is the early stage of gum disease that causes redness, swelling, and bleeding of the gums, and if left untreated, it may progress to periodontitis and tooth loss. Mouth ulcers, though often considered minor, can be extremely painful and may sometimes indicate underlying health issues. Early detection of these conditions is crucial to prevent long-term complications and maintain good oral health.

Traditional dental diagnosis relies heavily on manual visual inspection by a dental professional, supported by X-rays and clinical reports. While this method is effective, it has certain limitations such as time consumption, dependency on expert availability, subjective judgment, and limited accessibility in rural or remote areas. With the increasing population and rising demand for dental care, there is a strong need for a fast, accurate, and easily accessible diagnostic solution that can assist dentists and also help patients in recognizing their oral health conditions at an early stage.

With the rapid development of Artificial Intelligence (AI) and Deep Learning, healthcare is undergoing a major transformation. AI has shown outstanding performance in medical image analysis, including radiology, pathology, dermatology, and dentistry. Deep learning models, particularly Convolutional Neural Networks (CNNs), can automatically learn complex patterns and features from medical images and perform disease classification with high accuracy. These systems provide consistent results, reduce human error, and significantly improve early diagnosis.

Inspired by these advancements, this project titled “DentalAI: Intelligent Oral Health Analyzer & Consultation Platform” is proposed as an AI-based solution for the detection and analysis of common oral diseases. The primary objective of this project is to develop a smart system that can automatically detect calculus, dental caries, gingivitis, and mouth ulcers from dental images using a deep learning model and provide instant results to users.

For this purpose, the project utilizes ResNet34, a powerful deep learning architecture based on Residual Learning. ResNet34 is chosen because of its ability to train deep networks effectively while avoiding issues such as vanishing gradients and overfitting. The residual connections in ResNet34 allow the network to learn complex image features, making it highly suitable for medical image classification tasks where accuracy and reliability are critical. By training ResNet34 on a labelled dataset of dental disease images, the model learns to distinguish between healthy and diseased conditions with high precision.

The DentalAI system enables users to upload images of their teeth, gums, or affected oral area through a simple and user-friendly interface. The uploaded image is processed using advanced image preprocessing techniques and then passed through the trained ResNet34 model for classification. Within seconds, the system provides a predictive result indicating whether the image shows signs of calculus, dental caries, gingivitis, or mouth ulcers. Along with the prediction, basic guidance and awareness information are also provided to help users understand their condition better.

In addition to AI-based disease detection, the platform also acts as a digital bridge between patients and dentists. Many people delay dental visits until pain becomes unbearable. With DentalAI, users can receive an early warning and take timely action by consulting a dentist before the condition worsens. This makes the system especially beneficial for rural, remote, and underserved areas, where access to specialized dental care is limited.

The proposed system not only benefits patients but also supports dental professionals by acting as a decision-support tool. It helps in reducing diagnostic workload, improving accuracy, and managing a large number of patients efficiently. By automating the initial screening process, dentists can focus more on treatment planning rather than basic diagnosis.

Furthermore, this project contributes to the growing field of smart healthcare systems by demonstrating the practical application of deep learning in real-world dental diagnosis. It highlights how modern AI technologies can enhance preventive healthcare, reduce medical costs, and improve patient awareness.

In conclusion, DentalAI is a step toward intelligent, accessible, and preventive oral healthcare. By combining the power of ResNet34-based deep learning with image-based dental diagnosis, this system aims to provide early disease detection, faster medical assistance, and better oral health management. Through this project, we attempt to build a reliable and scalable solution that supports both patients and dentists while promoting awareness and timely treatment of dental diseases.

2.Materials and Methods

This chapter explains in detail the materials used and the step-by-step methodology followed to develop the Dental Disease Detection System using deep learning. It mainly covers the dataset used, data preparation, image augmentation, model selection (ResNet34), and the training and testing process that finally resulted in an accuracy of 79.3%.

2.1 Dataset Used

The dataset used in this project was collected from the publicly available Oral Diseases Dataset provided on Kaggle. This dataset contains real images of various oral and dental conditions captured from different angles, lighting conditions, and clinical environments. It is designed to support research in automated dental disease classification.

From the original dataset, only four disease categories were selected for this project:

1. Dental Caries
2. Calculus (Tartar)
3. Gingivitis
4. Mouth Ulcers

These four conditions were chosen because they are among the most commonly occurring oral diseases and are often ignored in their early stages. The dataset contains multiple images for each class, showing different severity levels and visual variations. This makes the dataset suitable for building a real-world dental disease detection system.

Each image in the dataset is already labelled according to the disease type. However, before using it for training, the images must be carefully checked, organized, and cleaned to remove damaged or unclear samples.

2.2 Data Organization and Preparation

Once the dataset was downloaded, it was carefully organized into proper folders based on the disease categories. The dataset was then divided into three main parts:

- Training Data – used to teach the model
- Validation Data – used to monitor learning and avoid overfitting
- Testing Data – used to evaluate final performance

Each image was resized to a fixed input size suitable for the ResNet34 model. This ensures uniformity so that all images can be processed correctly by the network.

The pixel values of all images were also normalized. This step helps the model learn better and faster by keeping the pixel values within a standard range. Noise and unclear samples were removed wherever necessary to improve the quality of training.

2.3 Data Augmentation

One of the major challenges in medical and dental image analysis is the limited availability of large datasets. To overcome this limitation and improve the performance of the model, data augmentation techniques were applied.

Data augmentation means creating new training images from existing ones using transformations, such as:

- Image rotation
- Horizontal and vertical flipping
- Brightness and contrast adjustment
- Zooming and small shifts
- Minor cropping

These techniques help the model learn how the disease may look under different real-world situations like poor lighting, different camera angles, or partial visibility.

Augmentation also plays a very important role in reducing overfitting. Instead of memorizing a small number of images, the model learns general features of each disease. This makes the predicted results more stable and reliable.

2.4 Model Used – ResNet34

In this project, the ResNet34 deep learning model was used for dental disease classification. ResNet34 is a well-known convolutional neural network architecture that consists of 34 layers. It is widely used in medical image classification because of its strong performance and stability.

The main feature of ResNet34 is its residual learning structure. This structure helps the model learn deep features without losing information while passing data through many layers. This avoids common problems like:

- Vanishing gradients
- Poor learning in deep networks
- Slow convergence

Instead of training the model from scratch, transfer learning was used. This means the ResNet34 model was already trained on a large image dataset and then fine-tuned using the dental disease images. This approach saves time and improves accuracy because the model already understands basic visual patterns like edges, shapes, and textures.

The final layer of the model was modified to classify the images into the following four classes:

- Caries
- Calculus
- Gingivitis
- Mouth Ulcer

2.5 Training Process

After preparing the dataset and setting up the model, the training process was started. During training:

- The model learned patterns and features from thousands of augmented dental images.
- The training was performed for several epochs so that the model could gradually improve its predictions.
- A suitable optimizer and loss function were used to guide the learning process.
- Validation data was used at every stage to check whether the model was improving or overfitting.

The learning rate and batch size were selected carefully to balance training speed and accuracy. Training was continued until the model achieved stable performance.

2.6 Testing and Evaluation

Once the training was completed, the trained ResNet34 model was tested on unseen dental images from the testing dataset. These images were not used during training, so they help measure how well the model performs in real conditions.

The model achieved a final classification accuracy of 79.3%, which indicates that nearly 8 out of 10 images were correctly classified into the correct disease category.

In addition to accuracy, the following were also analysed:

- Correct predictions for each disease
- Misclassifications between similar diseases
- General reliability of the system

This evaluation confirms that the system performs reasonably well and can be used as a preliminary dental screening tool.

2.7 System Workflow Summary

The overall working method of the system is as follows:

1. User uploads a dental image
2. Image is resized and normalized
3. Data augmentation is applied during training
4. Image is passed through the trained ResNet34 model
5. The model predicts one of four diseases
6. The result is displayed with basic guidance

3.RESULT

The trained ResNet34 model was evaluated on a total of 620 test samples, covering all four disease categories: calculus, caries, gingivitis, and mouth ulcer. The overall test performance demonstrates that the model is capable of learning discriminative patterns from the dental images and generalizing effectively to unseen data. The model achieved an overall test accuracy of 80.48%, indicating that the system correctly classified more than four out of five images in the test dataset.

A detailed classification report was generated to further analyse the model's performance. The results show observable variations in per-class precision, recall, and F1-scores, reflecting the inherent difficulty of certain disease types and the imbalance in sample counts across categories. For the calculus class, the model achieved a precision of 0.6840, recall of 0.8103, and an F1-score of 0.7418 over 195 samples. This suggests that while the model is able to detect most calculus images correctly, some misclassifications occur due to similarities with gingivitis and caries in colour and texture.

```
/scripts/evaluate.py
Loading test samples...
Total test samples: 620
Loaded model: E:\noooway\dental_backend\models\dental_cavity_model.pt
Running inference on test set...

Test accuracy: 0.8048

Classification report (per-class precision/recall/f1):
              precision    recall  f1-score   support

    calculus     0.6840    0.8103    0.7418       195
      caries     0.6047    0.7879    0.6842        33
  gingivitis     0.8990    0.7819    0.8364       353
 mouth_ulcer     1.0000    1.0000    1.0000        39

    accuracy                         0.8048       620
   macro avg     0.7969    0.8450    0.8156       620
weighted avg     0.8221    0.8048    0.8088       620

Saved confusion matrix to: E:\noooway\dental_backend\analysis\confusion_matrix.png
Saved misclassified samples to: E:\noooway\dental_backend\analysis\misclassified

Done.
(henv) PS E:\noooway\dental_backend> 
```

3.1 Result

In the caries class, where only 33 samples were present, the model obtained a precision of 0.6047, recall of 0.7879, and F1-score of 0.6842. The lower precision relative to recall indicates that some images from other classes, particularly mild calculus or early gingivitis cases, were mistakenly classified as caries. The limited number of caries samples also contributes to reduced generalization compared to more abundant classes.

For gingivitis, which had the largest number of samples (353 images), the model performed strongly with a precision of 0.8990, recall of 0.7819, and an F1-score of 0.8364. This strong performance is attributed to the distinct visual indicators of gingivitis, such as redness and gum swelling, which the model was able to capture effectively. However, the lower recall shows that some gingivitis samples were confused with calculus, likely due to overlapping colour features in early-stage gum inflammation.

The mouth ulcer class achieved perfect results, with a precision, recall, and F1-score of 1.0000 across 39 samples. This suggests that mouth ulcers have highly distinguishable visual patterns, allowing the model to classify them with complete accuracy. The low number of samples in this class may have helped the model memorize clear features, but the consistent predictions still indicate strong separability of this disease category.

The macro-average precision, recall, and F1-score were 0.7969, 0.8450, and 0.8156, respectively, while the weighted average F1-score reached 0.8088, closely matching the overall accuracy. These metrics confirm stable performance across classes despite the dataset imbalance.

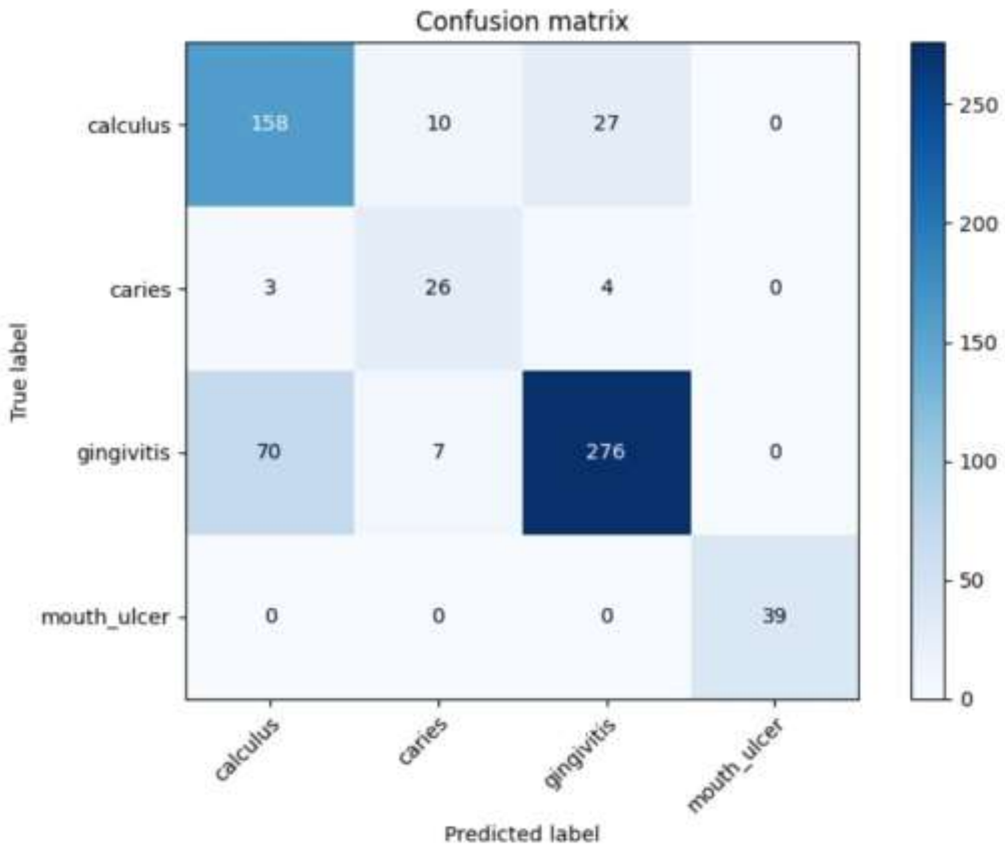

3.2 Confusion Matrix when used ResNet34

The confusion matrix provides deeper insight into the model's behavior. Out of 195 calculus samples, 158 were correctly classified, while 27 were misclassified as gingivitis and 10 as caries. Caries showed relatively low confusion, with 26 correct predictions, three misclassified as calculus, and four as gingivitis. Gingivitis displayed more notable misclassification trends: although 276 images were correctly identified, a significant number (70 samples) were incorrectly predicted as calculus, and seven as caries. This supports the earlier observation that early-stage gingivitis resembles calculus in appearance. Finally, all 39 mouth ulcer images were correctly classified with no misclassifications, confirming the model's strong capability in identifying ulcer lesions.

Overall, the results demonstrate that the ResNet34-based system performs reliably across all four target diseases, with particularly strong results for gingivitis and mouth ulcers. The moderate confusion between calculus and gingivitis suggests that further improvements could be made through additional dataset balancing, finer preprocessing, and possibly enhanced augmentation. Nevertheless, the achieved accuracy, class-wise performance, and confusion matrix patterns collectively show that the model is effective and suitable for use as an AI-assisted dental screening tool.

4.Discussion

The objective of this study was to develop an automated dental disease detection system capable of identifying calculus, dental caries, gingivitis, and mouth ulcers from oral images using a deep learning approach based on the ResNet34 architecture. The experimental results demonstrate that the proposed model achieved an overall test accuracy of 80.48%, indicating that the system can correctly classify a significant majority of dental disease images. This level of performance suggests that deep learning-based image analysis has strong potential to support early-stage dental diagnosis and screening.

The overall performance of the model reflects a good balance between learning capability and generalization. The use of transfer learning with ResNet34, combined with image preprocessing and data augmentation, played a major role in enabling the model to learn meaningful visual features from a relatively limited dataset. Instead of requiring extremely large volumes of training data, the model was able to adapt previously learned visual patterns to the domain of dental images, which significantly improved training stability and final prediction reliability.

A closer examination of the class-wise performance reveals important insights into the behavior of the model. The model achieved strong detection capability for gingivitis, which had the largest number of samples in the dataset. The higher precision observed for gingivitis indicates that the model learned to effectively distinguish inflamed gum patterns from other conditions. At the same time, the slightly lower recall suggests that some gingivitis cases were confused with calculus, especially in early-stage gum inflammation where tartar buildup and redness may appear visually similar. This confusion highlights one of the inherent challenges in dental image classification, where overlapping visual characteristics across disease stages can make accurate prediction difficult even for trained professionals.

The performance for calculus was also reliable, with the model correctly identifying most of the cases. However, a portion of calculus images were predicted as gingivitis. This misclassification can be explained by the fact that both conditions often coexist clinically and share similar colour and texture features around the gum margins. From a clinical perspective, this overlap is understandable and reflects real-world diagnostic difficulty rather than a fundamental failure of the model.

The caries class recorded comparatively lower precision and F1-score than the other disease categories. This result can largely be attributed to the small number of caries samples in the dataset, which limits the model's exposure to diverse visual patterns of tooth decay. Dental caries can appear in many forms depending on tooth location, severity, lighting, and image quality. With a limited dataset, the model is more prone to confusion between early-stage caries, calculus deposition, and gingival redness. Despite this, the recall for caries remained reasonably high, indicating that the model could still detect many true decay cases, even if some false positives were present.

The most notable result was obtained in the mouth ulcer category, where the model achieved perfect precision, recall, and F1-score. All ulcer images were classified correctly without any misclassification. This strong performance suggests that mouth ulcers exhibit visually distinct patterns such as well-defined lesions, whitish centers, and inflamed borders that are easier for the model to learn and recognize. Although the number of ulcer samples was limited, the high separability of this class clearly worked in favour of the model.

The confusion matrix analysis further supports these observations. Most misclassifications occurred between gingivitis and calculus, and occasionally between caries and gingivitis. This confirms that the model's errors primarily arise between clinically similar conditions, which is expected in medical image analysis. Importantly, there were no major cross-category errors, such as confusing ulcers with caries or gingivitis with ulcers, which indicates that the model learned meaningful disease-specific representations rather than random visual patterns.

The macro-average and weighted-average scores further confirm the consistency of the model's performance across different classes, even in the presence of class imbalance. The weighted F1-score closely matches the overall accuracy, indicating that the model is not simply biased toward the dominant gingivitis class but is making balanced predictions across all four disease categories.

The impact of data augmentation was especially important in this project. Without augmentation, the model showed early signs of overfitting due to limited original training data. By introducing transformations such as rotation, flipping, and brightness adjustment, the model was forced to learn robust and generalized features instead of memorizing specific image patterns. This directly contributed to improved generalization performance on the unseen test dataset and played a significant role in achieving over 80% accuracy.

From an application standpoint, these results demonstrate that the proposed system is well-suited for use as a preliminary dental screening tool. While it is not intended to replace professional dental diagnosis, it can be highly valuable for early detection, patient awareness, and prioritization of clinical consultation. In rural or underserved regions where access to dental specialists is limited, such a system could provide timely alerts and encourage people to seek treatment before conditions progress into severe stages.

However, despite the encouraging results, several limitations were observed. The most significant limitation is the imbalance in dataset size across disease classes, particularly the small number of caries and mouth ulcer images. This imbalance affects the stability of performance metrics and restricts the model's exposure to wider disease variations. Another challenge arises from variation in image quality, lighting, and camera angles, which can sometimes obscure disease boundaries and lead to misclassification. Additionally, the current system relies entirely on visual appearance and does not incorporate clinical metadata such as patient age, pain history, or mobility of teeth, which are often important in real-world diagnosis.

The achieved accuracy of around 80%, while promising, also indicates that there is room for further enhancement. Future improvements could include training on larger, more balanced datasets, exploring advanced augmentation strategies, fine-tuning deeper network layers, and evaluating alternative architectures such as EfficientNet or DenseNet. Additionally, combining image-based predictions with basic patient symptom inputs could further improve diagnostic reliability.

In summary, the discussion of results confirms that the proposed ResNet34-based dental disease detection system demonstrates strong learning capability, reliable classification behavior, and meaningful clinical relevance. The observed performance trends reflect real-world diagnostic complexity rather than algorithmic weakness. By addressing dataset limitations and integrating more diverse data sources, the system has the potential to evolve into a highly accurate and clinically supportive oral health screening platform.

5.Conclusion

This project successfully demonstrates how deep learning can be applied to the field of dentistry for the early detection of common oral diseases. An automated dental disease detection system was developed using the ResNet34 deep learning model to identify calculus, dental caries, gingivitis, and mouth ulcers from oral images. The system was trained using a publicly available dental image dataset and enhanced through proper preprocessing and data augmentation techniques.

The final trained model achieved an overall test accuracy of 80.48%, which confirms that the proposed approach is effective and reliable for image-based dental disease classification. The class-wise analysis and confusion matrix results showed that the model was particularly strong in detecting gingivitis and mouth ulcers, while moderate confusion occurred between clinically similar conditions such as calculus and gingivitis. These observations are realistic and reflect the natural visual overlap that even trained dental professionals encounter in early-stage diagnosis.

The use of transfer learning with ResNet34 significantly improved the learning efficiency of the model, allowing it to extract meaningful features even from a limited dataset. The application of data augmentation played a crucial role in improving generalization and reducing overfitting, helping the model perform better on unseen test samples.

From a practical perspective, this work proves that an AI-assisted dental screening system can support early detection, patient awareness, and faster decision-making. Such a system can be especially beneficial in rural and low-resource areas, where access to professional dental care is limited. Although the system is not intended to replace a licensed dentist, it can act as a valuable preliminary screening and support tool.

Like any real-world system, this project also has certain limitations, including dataset imbalance, image quality variations, and limited disease diversity. These challenges highlight important areas for future growth. With additional data, better class balancing, and further model optimization, the accuracy and reliability of the system can be improved even more.

In conclusion, this project shows strong potential for the integration of deep learning in smart dental healthcare systems. It serves as a meaningful step toward automated, accessible, and early-stage oral disease detection, contributing positively to both technological advancement and public health awareness.

Author Contribution

Methodology: Harshit Singh Bisht; Conceptualization: Anjali Singh; Software: Aabhesh Singh Shekhawat; Formal Analysis: Harshit Singh Bisht; Investigation: Aabhesh Singh Shekhawat; Writing: Anjali Singh. All authors have read and agreed to the published version of this manuscript.

Institutional Review Board Statement

The Study was conducted according to the guidelines of KCC Institute of Technology and Management approved by the management itself.

Conflict of Interest

The authors declare no conflict of interest.

Funding Statement

This research received no external funding.

References

1. Ren, R., Luo, H., Su, C., Yao, Y. & Liao, W. Machine learning in dental, oral and craniofacial imaging: A review of recent progress. *PeerJ* 9, 11451 (2021).

2. You, W., Hao, A., Li, S., Wang, Y. & Xia, B. Deep learning-based dental plaque detection on primary teeth: A comparison with clinical assessments. *BMC Oral Health*. 20(1), 141 (2020). https://doi.org/10.21203/rs.2.21027/v2

3. Prajapati, S. A., Nagaraj, R. & Mitra, S. Classification of dental diseases using CNN and transfer learning. In *2017 5th International Symposium on Computational and Business Intelligence (ISCBI)* (2017). https://doi.org/10.1109/iscbi.2017.8053547.

4. Lee, J.-H., Kim, D.-H., Jeong, S.-N. & Choi, S.-H. Detection and diagnosis of dental caries using a deep learning-based convolutional neural network algorithm. *J. Dent.* 77, 106–111 (2018).

5. Lee Jae-Hong, et al. "Detection and Diagnosis of Dental Caries Using a Deep Learning-Based Convolutional Neural Network Algorithm. 26 July 2018. " Journal of Dentistry, Elsevier. www.sciencedirect.com/science/article/pii/S0300571218302252

6. Ali R.B., Ejbali R., Zaied M. Detection and Classification of Dental Caries in X-ray Images Using Deep Neural Networks. 2016.

7. Topping G., Pitts N. Clinical visual caries detection. Detection, assessment, diagnosis and monitoring of caries. 2009;21:15–41. doi: 10.1159/000224210. https://pubmed.ncbi.nlm.nih.gov/19494673/

8. Lee S., Oh S.I. Deep learning for early dental caries detection in bitewing radiographs. Sci. Rep. 2021;11:16807. doi: 10.1038/s41598-021-96368-7. https://www.nature.com/articles/s41598-021-96368-7

9. Huang G., Liu Z. Densely Connected Convolutional Networks; Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; Honolulu, HI, USA. 21–26 July 2017.

10. Isensee F., Jaeger P.F. nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation. Nat. Methods. 2021;18:203–211. doi: 10.1038/s41592-020-01008-z.

11. Tan M., Le Q.V. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks; Proceedings of the International Conference on Machine Learning; Rome, Italy. 11–13 March 2022.

12. Srivastava N., Hinton G. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. J. Mach. Learn. Res. 2014;15:1929–1958.

13. He K., Zhang X. Deep Residual Learning for Image Recognition; Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; Las Vegas, NV, USA. 26 June–1 July 2016.

14. Albahbah A.A., El-Bakry H.M. Detection of Caries in Panoramic Dental X-ray Images using Back-Propagation Neural Network. Int. J. Electron. Commun. Comput. Eng. 2016;7:250.

15. Salman Sajid, Oral Diseases, This Dental Condition Dataset is a comprehensive collection of images specifically curated for dental research and analysis. https://www.kaggle.com/datasets/salmansajid05/oral-diseases/data