AI-Based Dental Disease Detection Using OPG Images

Mrs. L. Velvizhi[1], S. Sujatha[2]

[1]*Assistant Professor, Department of Computer Science and Engineering, Vivekanandha College of Engineering for Women, Tamilnadu, India*

[2]*PG Scholar, Department of Computer Science and Engineering, Vivekanandha College of Engineering for Women, Tamilnadu, India*

[1]Email: velvizhi@vcew.ac.in

[2]Email: suji20232024@gmail.com

Abstract— Panoramic dental radiographs, known as orthopantomograms (OPGs), give a wide view of the teeth and surrounding jaw structures. Interpreting them can be complicated by image quality, anatomical overlap and variation between imaging devices. This paper proposes an AI-based workflow for the preliminary analysis of OPG images using a convolutional neural network (CNN) or a transfer-learning model. The planned system comprises image validation, preprocessing, model inference and a simple mobile interface through which a user can upload an image and view the result. Dataset selection, annotation quality, patient privacy and class-wise evaluation are treated as core design requirements, and performance will be assessed with accuracy, precision, recall, F1-score and a confusion matrix rather than accuracy alone. The work is currently at the design stage, so no trained-model performance or clinical results are claimed. The intended outcome is a research prototype that can support preliminary review while leaving diagnostic decisions to qualified dental professionals.

Keywords— Artificial intelligence, dental imaging, OPG, panoramic radiograph, convolutional neural network, transfer learning, image classification, mobile application.

I. INTRODUCTION

Dental radiography is a key source of information for examining teeth, roots and the surrounding jaw. An orthopantomogram (OPG) records both jaws in a single panoramic image, helping clinicians assess tooth presence and position, impacted teeth and certain bone or jaw findings. Panoramic images are, however, two-dimensional projections. Positioning errors, magnification, blur, overlapping structures and differences between imaging devices can all reduce how clearly a feature appears.

Oral diseases remain a substantial global health burden [6], which has encouraged interest in tools that support dental image review. Artificial intelligence (AI), and deep learning in particular, has been widely studied for image-recognition tasks in healthcare. Convolutional neural networks learn visual features directly from image data and can be adapted to a new task through transfer learning [1], [2]. In dental imaging, how useful such a system is depends on the condition under study, the quality of the reference labels and how representative the images are. A model trained for one narrow task should not be assumed to recognize every dental disease.

This work proposes a structured pipeline for an OPG image-analysis prototype. Through a mobile interface, the user selects an image; the system then validates and preprocesses it, passes it to a trained model and displays a preliminary output. The project is intended as a research demonstration, not as a substitute for examination by a dentist.

A. Problem Statement

Interpreting OPG images manually demands clinical knowledge and close attention to image quality and anatomical context. A computer-based screening aid could help apply image processing consistently and draw attention to selected findings. However, an incorrect prediction could mislead users if the system presents its output as a confirmed diagnosis. The project therefore requires a clearly defined classification task, reliable labels, independent evaluation and an explicit safety notice.

B. Objectives and Contributions

The project aims to design and evaluate a prototype that detects selected dental conditions in OPG images. The planned objectives are to review the relevant literature, identify a permitted and annotated dataset, choose disease categories that the labels support, implement preprocessing, build a CNN or transfer-learning baseline, evaluate class-wise performance and design a mobile interface for submitting images and viewing results. The main contribution is an end-to-end plan that links data handling, model evaluation and user-facing output without claiming clinical performance before testing.

C. Scope

The current phase covers the problem definition, technical background, dataset requirements, proposed model, evaluation plan and application design. Implementation and experiments are planned for the next phase, subject to dataset access and supervisor approval. The final model will report only the categories represented in the selected dataset. No accuracy values or clinical outcomes are reported at this stage.

II. RELATED WORK AND TECHNICAL BACKGROUND

A. AI in Medical and Dental Image Analysis

Deep learning is now an established research approach in medical image analysis, covering classification, detection and segmentation tasks [1], and its wider role in medicine has been widely discussed [8]. CNNs apply convolutional operations to learn local patterns and combine them into higher-level representations [3]. In dentistry, the target task must be defined carefully, because tooth detection, tooth numbering, anatomical segmentation and disease classification are distinct problems. Success on one of them does not demonstrate performance on another.

Research on AI in dentistry highlights its potential for decision support, together with concerns about data quality, explainability, external validity and accountability for clinical decisions [7]. A model may learn patterns tied to the source of an image rather than to the intended condition, especially when the training data are small or come from only a few devices. Dataset documentation and validation are therefore as important as the choice of network architecture.

B. CNNs and Transfer Learning

A CNN typically stacks convolutional layers, nonlinear activations and pooling or strided operations to extract image features, and a classification head then maps the learned representation to scores for the target classes [3]. Transfer learning starts from parameters learned on a larger image dataset and adapts them to the new task. This can offer a practical baseline when the available dental dataset is small, although its benefit for medical images must be confirmed experimentally [5]. Controlled training and careful evaluation are needed to limit overfitting [2], [4].

C. Dataset Quality and Evaluation

Images should be taken only from sources that permit the intended research use. If institutional images are considered, the necessary permissions and ethics approval must be secured before use. Personal

identifiers should be removed wherever possible, and patient information should not appear in the application output. The annotation process should define each target class and record how uncertain or poor-quality images are handled.

Accuracy alone can hide weak performance on uncommon classes. Precision is the proportion of positive predictions that are correct, recall is the proportion of reference-positive cases that are identified, and the F1-score balances the two. A confusion matrix shows which classes are being mistaken for one another. Where feasible, these metrics should be reported for each supported class, together with the dataset size and class distribution.

D. Research Gap

A useful prototype must address more than the construction of a neural network. Practical challenges include limited or imbalanced data, inconsistent labels, variation in image acquisition, data leakage between training and testing, and unclear communication of prediction uncertainty. At the design level, this project responds by defining a limited target scope, documenting preprocessing, reserving a held-out test set, reporting several metrics and presenting results as preliminary. These safeguards reduce avoidable risks but do not establish clinical validity.

III. PROPOSED SYSTEM

The proposed system has five functional modules: image input, preprocessing, AI inference, result generation and a mobile user interface. The image-input module accepts a supported OPG file and checks that it can be decoded. The preprocessing module applies the same resizing, intensity handling and normalization steps that were used during model training. The inference module loads a trained CNN-based model and produces scores for the selected classes. The result module converts those scores into a structured response, and the mobile interface displays the output together with a clear limitation notice.

TABLE I. PROPOSED SYSTEM MODULES

Module	Input	Responsibility
Image input	OPG file	Select a file and check its format and readability
Preprocessing	Validated image	Resize and normalize consistently
AI inference	Processed image	Generate scores for supported classes
Result generation	Class scores	Format the preliminary output and handle errors
Mobile interface	User action and result	Upload the image and display the result with a safety notice

A. System Workflow

The planned workflow is as follows: (1) the user opens the application and reads the intended-use notice; (2) an OPG image is selected from the device; (3) the system checks the file format and image validity; (4) the image is preprocessed; (5) the model generates class scores; and (6) the application displays the supported output with a reminder to consult a qualified dental professional. Unsupported files, corrupted images, connection failures and inference errors should trigger clear error messages rather than misleading results.

B. Dataset Plan

The dataset will be chosen after reviewing its access conditions, license terms, image type, annotation method and suitability for the target classes. Where permitted, each record should include a study identifier, image path, reference label, source information and an image-quality note. If several images belong to the same patient, they should be kept in the same data split to reduce the risk of information leakage. The training, validation and test sets will be defined before model selection, and the test set will remain untouched until the final evaluation.

C. Image Preprocessing

Preprocessing will standardize the input while preserving clinically relevant detail. The initial steps include file validation, a check of image dimensions, conversion to the representation the model expects, resizing and normalization. Training augmentation may include modest rotations, translations, zooming or brightness changes, but only where these preserve the anatomical meaning of the image. Augmentation will not be applied to validation or test images. The exact preprocessing pipeline will be documented so that it can be reproduced.

D. Proposed Model

A CNN-based classifier will be evaluated as the primary model family. A transfer-learning model may be used to initialize the feature extractor, after which the final classification layers will be adapted to the selected categories. The model will output a score for each supported class. The final architecture and training settings will depend on the dataset size, the available computing resources and the observed validation performance. Model confidence will not be presented as proof that a condition is present.

IV. METHODOLOGY

A. Training Strategy

The dataset will be divided into training, validation and test subsets before training begins. The training set will be used to update model parameters, the validation set will guide architecture and hyperparameter choices, and the held-out test set will be reserved for final reporting. If patient identifiers or source groups are available, the split will be designed so that related images do not end up in different subsets. The learning rate, batch size, number of epochs and regularization settings will be chosen through controlled experiments, and early stopping may be applied once validation performance stops improving.

B. Evaluation Metrics

The planned evaluation will include accuracy, precision, recall, F1-score and a confusion matrix. Class-wise values will be reported where possible, because an overall average can conceal poor detection of a less frequent category. Macro-averaged scores may be used when every class is given equal importance, and weighted scores may be added to reflect class frequency. The report will also describe the data split and any exclusions alongside the metrics.

TABLE II. PLANNED EVALUATION MEASURES

Measure	Purpose
Accuracy	Overall proportion of correct predictions
Precision	Proportion of positive predictions that are correct
Recall	Proportion of reference-positive cases that are identified
F1-score	Balance between precision and recall
Confusion matrix	Correct and incorrect predictions for each class

C. Mobile Application Design

The mobile application is planned around a short, clear sequence: open the app, select an OPG image, submit it for analysis and review the returned output. A Flutter-based front end may communicate with a lightweight backend API that validates the upload, applies the saved preprocessing pipeline, loads the model and returns a structured response. The interface should flag unsupported input, avoid showing categories that the model was not trained to predict, and state that the output is experimental and not a medical diagnosis.

D. Privacy, Safety and Ethics

Dental radiographs can contain sensitive health information. Only data that are ethically and legally available for the stated research purpose will be used. Personal identifiers should be removed wherever possible, files should be kept in restricted locations and access should be limited to authorized project members. In line with general AI risk-management guidance [9], the system's limitations will be documented and communicated to users. Because it may produce false-positive and false-negative predictions, it must not recommend treatment or make definitive diagnostic claims, and users should be directed to a qualified dental professional for interpretation.

V. PLANNED EVALUATION AND DISCUSSION

Implementation is planned in stages: confirming the target classes and dataset permissions, preparing and checking the images, establishing a reproducible training pipeline, training a baseline model, evaluating it on held-out data, connecting inference to the mobile interface and testing the complete workflow. Application tests will cover a valid image, an unsupported file, a corrupted image, an image with very small dimensions, a network interruption and a server-side inference error.

The analysis will begin by describing the dataset, including the image count, class distribution, source and exclusions. Training and validation curves will be examined for signs of overfitting. The final model will be assessed with the planned metrics, and errors will be reviewed to identify likely causes such as image quality, label ambiguity, limited examples or overlapping radiographic appearance. Any conclusions will be limited to the tested dataset and task.

No measured results are available because the prototype and experiments have not yet been completed. Performance values will be reported only after training and evaluation. Even if results prove promising on a curated dataset, further testing on independent data and clinical review would be required before any real-world use is considered.

VI. CONCLUSION AND FUTURE WORK

This paper has presented a proposed AI-based workflow for the preliminary analysis of dental OPG images. The design links image validation, preprocessing, CNN-based inference, result presentation and a mobile interface. It also identifies dataset governance, label quality, class-wise evaluation and cautious communication of model output as essential parts of the work. The prototype is intended to support research and demonstration, not to replace clinical judgment.

Future work will identify a permitted and suitably annotated OPG dataset, finalize the target classes, implement and train a baseline model, and evaluate it on a held-out test set. The trained model will then be integrated into the mobile application and tested for input validation, error handling and usability. Any future performance claims will rest on measured results and be accompanied by clearly stated limitations.

REFERENCES

[1] G. Litjens et al., "A survey on deep learning in medical image analysis," Medical Image Analysis, vol. 42, pp. 60–88, 2017.

[2] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, pp. 436–444, 2015.

[3] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.

[4] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," Journal of Big Data, vol. 6, article 60, 2019.

[5] M. Raghu, C. Zhang, J. Kleinberg, and S. Bengio, "Transfusion: Understanding transfer learning for medical imaging," in Advances in Neural Information Processing Systems, vol. 32, 2019.

[6] World Health Organization, Global Oral Health Status Report: Towards Universal Health Coverage for Oral Health by 2030. Geneva, Switzerland: WHO, 2022.

[7] F. Schwendicke, W. Samek, and J. Krois, "Artificial intelligence in dentistry: Chances and challenges," Journal of Dental Research, vol. 99, no. 7, pp. 769–774, 2020.

[8] E. J. Topol, "High-performance medicine: The convergence of human and artificial intelligence," Nature Medicine, vol. 25, pp. 44–56, 2019.

[9] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, 2023.