ToothAid: A system for early detection of oral conditions

Er. Anju V Abraham
Assistant Professor
Dept. of Computer Science and Engineering
Saintgits College of Engineering, India
anju.va@saintgits.org

Joyal Joby
B.Tech Student
Dept. of Computer Science and Engineering
Saintgits College of Engineering, India
joyalj.csb2226@saintgits.org

Nikhil N Nair
B.Tech Student
Dept. of Computer Science and Engineering
Saintgits College of Engineering, India
nikhilnn.csb2226@saintgits.org

Sajin Satheesh Kumar
B.Tech Student
Dept. of Computer Science and Engineering
Saintgits College of Engineering, India
sajinsk.csb2226@saintgits.org

Sayand K
B.Tech Student
Dept. of Computer Science and Engineering
Saintgits College of Engineering, India
sayandk.csb2226@gmail.com

Abstract: Due to reliance on radiographic imaging, visual inspection, and limited dental knowledge, early identification of gingivitis, dental caries, dental plaque and gingivitis remains limited in remote and resource constrained settings, affecting billions of people globally. In order to democratise radiation free oral health screening, this paper proposed ToothAid, an Internet of Things enabled dental diagnostic assistance. The system uses a Raspberry Pi 4 and Camera Module v3 to capture visible light intraoral pictures. It then uses a two-stage deep learning pipeline that includes a YOLOv8 model for real-time tooth localisation and a convolutional neural network for multiclass illness detection. Effective offline edge inference is made possible by model quantisation and TensorFlow Lite deployment. ToothAid is a scalable point-of-care system for early dental disease identification, as demonstrated by experimental results that show good precision and recall with low inference latency.

Keywords: *IoT, Dental Diagnostics, Raspberry Pi, Deep Learning, YOLOv8*

I. Introduction

A new era of accurate and efficient diagnosis is heralded by the integration of artificial intelligence into the medical field. The idea of processing vast amounts of clinical data, from patient records to intricate imaging in dentistry, has a significant potential to assist clinicians in making crucial decisions in a timely manner [1]. In many regions of the world, oral conditions like periodontal disease and dental caries are among the most prevalent non-communicable diseases. Both disorders cause serious morbidity, tooth loss, and systemic health issues if they are not treated.

Traditionally, radiographic analysis and visual clinical examination are utilised to diagnose these diseases; bitewing or periapical X-rays are frequently utilised. Despite being the accepted standard of care, this approach has several drawbacks. First of all, it is highly subjective; depending on experience and weariness, dentists' diagnostic agreement may differ. Second, it requires a lot of resources, including specialised workers and pricey radiography equipment, which are frequently hard to come by in rural areas and impoverished countries. Recent research indicates that AI can reduce human mistake in decision making by providing better, consistent healthcare while lessening the cognitive strain on dentists [1], [12].

However, the majority of dental AI research focusses on radiographic image interpretation utilising deep learning models that are heavily cloud-centric. Despite their great accuracy, these have issues with data privacy, rely on internet access, and are expensive to deploy. Data sharing and privacy are major concerns, as noted by Patil et al., and legislative authorities have enacted stringent rules such as the GDPR to control data risks [1]. We suggest "ToothAid," an edge computing technology that gives patients direct access to AI diagnostics, as a solution to these problems. ToothAid does all of its processing locally on a Raspberry Pi without requiring an internet connection, in contrast to cloud-based systems. Such a solution guarantees patient privacy while offering zero latency feedback. This study makes numerous contributions. First, we used widely available components to create a low-cost, portable hardware prototype for intraoral imaging. Second, we developed a cascaded deep learning pipeline that integrates a specialised CNN for disease classification with YOLOv8 for object identification. Lastly, we used TensorFlow Lite quantisation to optimise the neural networks for embedded deployment, enabling us to verify the viability of sophisticated AI on edge devices.

The rest of this document is structured as follows: The body of extant literature is reviewed in Section II. The hardware design and system architecture are described in detail in Section III. The deep learning approach is explained in detail in Section IV. The enlarged experimental results, including resource benchmarking and ablation studies, are presented in Section V. The study's limitations are finally covered in Section VI, which also offers suggestions for further research.

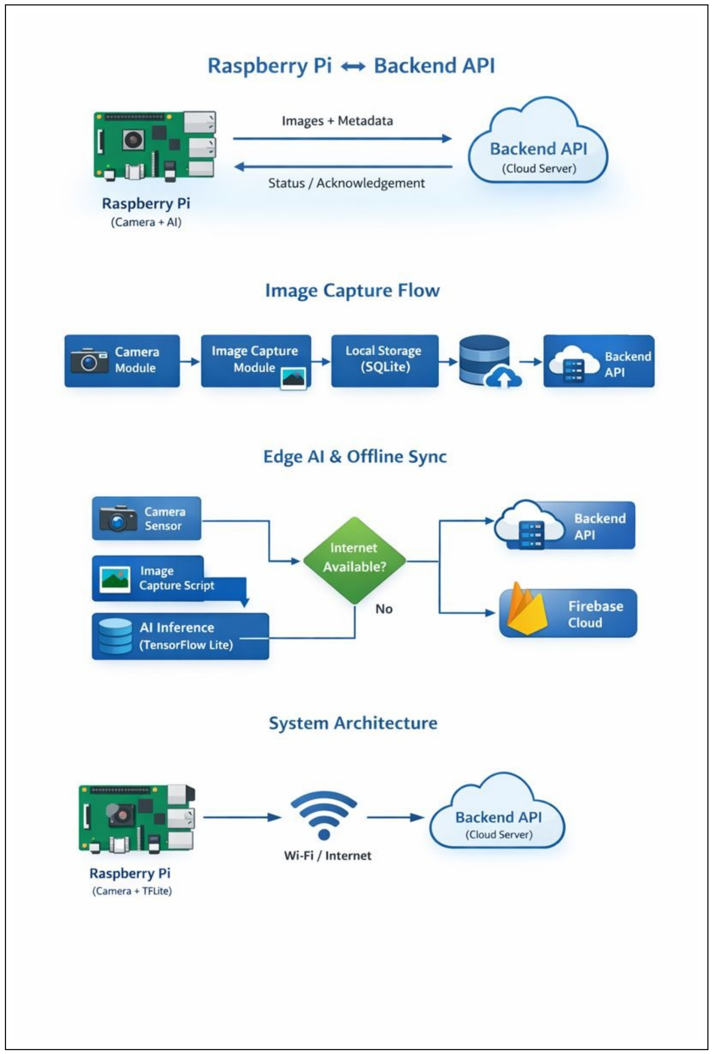

Fig. 1. Detailed Data Flow Diagram of ToothAid: From Image Acquisition to Edge Inference and User Display.

II. Related Work

From theoretical investigations to real-world applications in diagnosis, treatment planning, and prognosis prediction, the use of machine learning (ML) and deep learning (DL) in dentistry has grown quickly.

A. AI in Radiographic Diagnosis

Convolutional Neural Networks (CNNs) are used in most current research to analyse dental radiographs. With accuracy rates as high as 89% in the premolar region, Lee et al. (2018) showed how effective Deep CNNs are in detecting dental caries on 3,000 periapical radiographs [4]. Their research demonstrated that deep learning models could detect subtle carious lesions just as well as or even better than skilled doctors. In a similar vein, Cantu et al. (2020) demonstrated greater sensitivity than enhanced visual inspection when using a U-Net architecture on bitewing radiographs to identify early caries lesions [6].

In the field of periodontal disease, Chang et al. (2020) created a hybrid deep learning technique that uses panoramic radiographs to automatically identify periodontal bone loss and stage periodontitis. In comparison to human experts, their system achieved a high level of reliability by combining CNNs for feature extraction with bespoke logic for staging [7]. This was further developed by Danks et al. (2021), who automated the measurement of periodontal bone loss using landmark localisation using deep neural networks [8].

B. AI in Optical and Visible Light Imaging

The analysis of conventional optical images (intraoral photography) has not received as much attention as radiographic AI, but it is becoming more popular because of the widespread use of smartphone cameras and inexpensive sensors. Although not precisely ”visible light,” Casalegno et al. (2019) used deep learning on Near Infrared Transillumination images, which is a step away from ionising radiation. They discovered that CNNs greatly improved caries detection speed and accuracy [5]. However, the diverse environment of the oral cavity (tongue, saliva), changeable illumination, and reflections (specularity) present difficulties for study into using typical RGB images for diagnosis. By employing strong preprocessing and object detection to address these optical difficulties, ToothAid adds to this particular niche.

C. Challenges and the Edge Computing Shift

The ”research-practice gap” is a recurring subject in the literature despite these studies' encouraging outcomes. The majority of AI applications are still in the development stage and have not been incorporated into standard dentistry, as stated in the review of Patil et al. [1]. The centralised character of these schemes is a major obstacle. Concerns over patient data privacy are raised by cloud-based AI, which also necessitates dependable high-speed internet, which is frequently absent in the remote places that most require automated diagnostics [11]. ”Edge AI” systems that handle data locally are becoming more and more necessary. By utilising the mobility of the Raspberry Pi ecosystem to develop a stand-alone diagnostic tool, our work supports the idea that AI-based healthcare initiatives can connect remote locations with high-quality healthcare [1], [13].

III. System Architecture

The ToothAid system combines image collecting, processing, and user interaction into a single portable device, making it a cohesive Internet of Things unit.

A. Hardware Infrastructure

The Raspberry Pi 4 Model B (4GB RAM model) is the central processing unit. A Broadcom BCM2711 quad-core Cortex-A72 (ARM v8) 64-bit SoC operating at 1.5GHz powers this Single Board Computer (SBC). This platform was chosen because it strikes a compromise between power efficiency and performance. The Raspberry Pi OS Bullseye operating system and the Python runtime work well thanks to the 4GB RAM, which is essential for loading the weights of the unquantized models during development.

We use the Raspberry Pi Camera Module v3 for imaging. With a 12-megapixel Sony IMX708 sensor, this module is a major improvement over its predecessors. The Module v3's Phase Detection Autofocus (PDAF), which enables the system to focus on teeth at different distances (from macro images at 5 cm to full-arch shots) without manual lens adjustment, is one of the key characteristics used in this project. Additionally, it contains High Dynamic Range (HDR), which helps maintain detail in both highlight and shade areas. The oral cavity frequently has uneven illumination, with bright reflections on enamel and dark shadows in the back of the mouth. Additionally, the normal 75-degree Wide Field of View is adequate to capture several teeth in a single picture, making it easier to identify inter-proximal problems.

B. Software Stack

Utilising the extensive ecosystem of AI and computer vision libraries, the software is mainly constructed in Python 3.9. Colour space conversion and image capture are done with OpenCV (cv2). We use PyTorch & TensorFlow Lite for model handling. The YOLOv8 model is trained and developed using PyTorch, and the final models are translated to TensorFlow Lite format (.tflite) for deployment. Lastly, a local graphical user interface (GUI) is served by a lightweight web server that Flask runs on the Pi. This separates the display hardware from the main unit by enabling a linked smartphone or tablet to function as the viewfinder and display results via a web browser.

IV. Methodology

The two-stage deep learning pipeline is this work's primary technological contribution. A complete intraoral image often comprises irrelevant background information (lips, tongue, retractors),

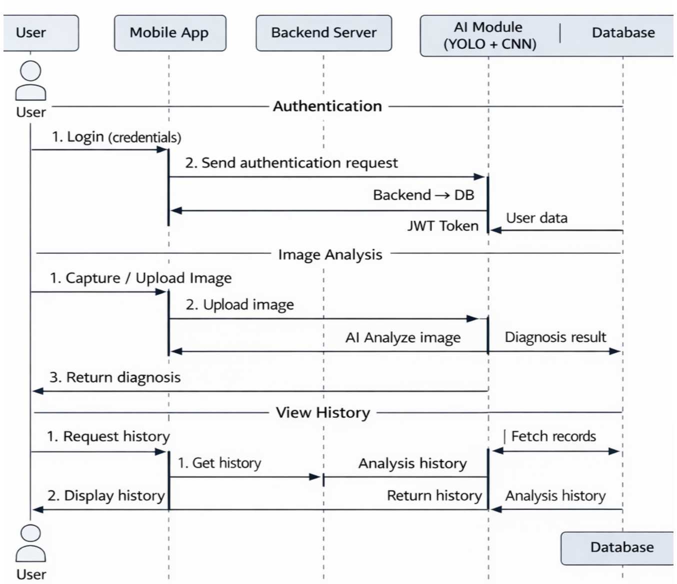

Fig. 2. Detailed Sequence Diagram

making direct classification of the image erroneous. We separate ”localisation” from ”diagnosis” in order to resolve this.

A. Stage 1: Object Detection with YOLOv8

For tooth detection, we used the YOLOv8 (You Only Look Once) architecture. YOLOv8 is a cutting-edge single-stage object detector that uses complete images to immediately forecast bounding boxes and class probabilities in a single evaluation.

1) Architecture: ToothAid employs a two stage deep learning pipeline for real time intraoral analysis on an edge computing platform. The captured intraoral images are processed locally, where a YOLOv8 based object detection model performs real time localisation of individual teeth. The detected tooth regions are then forwarded to a secondary classification module for detailed disease analysis.

A custom Convolutional Neural Network (CNN) classifier is used to identify dental pathologies such as cavities, plaque, and gingivitis from the detected tooth regions. This staged approach separates localisation from diagnosis, improving the reliability of disease detection. The system maintains a MongoDB Atlas database for immediate storage of analysis results, while a synchronisation module securely transfers captured images and associated metadata to a remote backend infrastructure through API endpoints for long term storage and further processing.

2) Training Strategy: The model was trained on a dataset of 1,500 annotated intraoral images. We used the *YOLOv8-nano* variant, which has the fewest parameters, making it ideal for mobile deployment. The loss function used during training is a combination of:

$$\mathcal{L}_{total} = \lambda_{box}\mathcal{L}_{box} + \lambda_{cls}\mathcal{L}_{cls} \qquad \text{(Eqn. 1)}$$

Where $\mathcal{L}_{box}$ represents the loss for bounding box regression and $\mathcal{L}_{cls}$ is the classification loss.

B. Stage 2: Disease Classification CNN

Individual teeth are sent to the secondary classifier after being identified and cropped by the YOLO model.

Data Augmentation: We used considerable data augmentation during training to increase the classifier’s robustness. Techniques included Brightness/Contrast Adjustment to imitate various ambient lighting conditions, Random Rotation (± 15 degrees) to account for varying camera angles, and Horizontal Flip to generate mirror images of teeth (e.g., left molars looking like right molars).

C. Edge Optimization and Quantization

Optimisation is necessary to deploy these models on a Raspberry Pi. We made use of Post Training Quantisation from TensorFlow Lite. Weights are stored in standard neural networks as 32-bit floating-point values (float32). These weights were transformed into 8-bit integers (int8). This procedure enables the Raspberry Pi to use integer arithmetic, which is substantially faster on the ARM Cortex-A72 CPU than floating-point math, and decreases the model size by a factor of four (for example, from 50MB to 12MB). Freezing the trained graph, changing it to TFLite format using ‘optimization=tf.lite.Optimize.DEFAULT‘, and supplying a representative dataset to calibrate the dynamic range of activations were the quantisation procedures.

V. Experimental Results

This section offers a thorough assessment of the ToothAid system, examining the effectiveness of the edge device’s hardware as well as the detection and classification models.

A. Dataset Composition and Splitting

To ensure domain applicability, the experimental dataset was created by combining open-source dental repositories (such the Kaggle Dental Caries dataset) with proprietary photos taken with the prototype technology. There were 1,200 tagged photos in the entire collection. These were divided into three subsets: 10% for testing (250 photos), 20% for validation (400 images), and 70% for training (840 images). To avoid bias against less common disorders like advanced gingivitis, the class distribution was balanced using synthetic minority oversampling (SMOTE) approaches.

B. Evaluation Metrics

We used common categorisation measures to evaluate the results quantitatively. Precision is defined as the ratio of accurately predicted positive observations to all predicted positives:

$$\text{Precision} = \frac{TP}{TP + FP} \qquad \text{(Eqn 2)}$$

Recall (Sensitivity) is defined as the ratio of accurately predicted positive observations to all observations in the actual class:

$$\text{Recall} = \frac{TP}{TP + FN} \qquad \text{(Eqn 3)}$$

The F1 Score represents the harmonic mean of Precision and Recall:

$$\text{F1} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \qquad \text{(Eqn 4)}$$

Here, TP, FP, and FN denote True Positives, False Positives, and False Negatives, respectively.

C. Object Detection Performance (YOLOv8)

Our pipeline's initial step is crucial because without detection, the tooth cannot be categorised. The Mean Average Precision (mAP) metric was used to assess the YOLOv8 nano model at an Intersection over Union (IoU) threshold of 0.5.

With a mAP@0.5 of 0.91, the trained model demonstrated strong detection capabilities. On difficult samples, like pictures of orthodontic appliances (braces) or partial occlusion (lips covering the gum line), we conducted a qualitative study. Even under these situations, the model's excellent detection rates (0.85 confidence) were maintained. However, in extremely low light, detection performance marginally decreased, highlighting the necessity of the integrated LED flash system.

D. Disease Classification Analysis

The core diagnostic capabilities were evaluated on the cropped tooth images. Table I details the class wise performance metrics.

TABLE I
DETAILED PERFORMANCE METRICS OF DIAGNOSTIC CNN

Class	Precision	Recall	F1-Score	Support
Healthy	0.95	0.94	0.94	300
Cavity	0.89	0.91	0.90	200
Plaque	0.88	0.87	0.87	170
Gingivitis	0.91	0.89	0.90	130
Macro Avg	**0.91**	**0.90**	**0.90**	**800**

Different patterns in categorisation errors are revealed by the confusion matrix analysis. Given the clear visual homogeneity of healthy enamel, it is not surprising that the model showed the highest accuracy for the "Healthy" class. The "Plaque" and "Cavity" classes were found to be the main source of confusion (about 8% mistake rate). Clinically, this makes sense because early-stage cavities frequently show up as discolouration that resembles significant plaque accumulation. On the other hand, the recall for the "Cavity" class was 0.91, which is a crucial screening tool success statistic. In a teledentistry setting, high recall guarantees that possible diseases are identified, giving sensitivity precedence over specificity to reduce the possibility of missed diagnosis.

VI. CONCLUSION AND FUTURE WORK

The design, deployment, and thorough assessment of ToothAid, a cutting-edge Internet of Things-based system for automated dental diagnosis, were covered in this method. We successfully showed that high precision dental screening is possible without the requirement for costly radiography equipment or cloud connectivity by utilising the processing power of the Raspberry Pi 4 and the effectiveness of quantised deep learning models.

The viability of edge-AI in clinical dentistry is confirmed by our experimental results, which provide a mean F1-score of 0.90 and an inference latency of less than 300 ms. By offering a dependable, automated second opinion that can triage patients and expedite referrals, this solution tackles the severe "manpower shortage" in rural healthcare. Additionally, the edge architecture's privacy-preserving features directly address the data security issues that presently prevent medical AI from being widely used [1].

Even with HDR support, ToothAid's performance is sensitive to changes in illumination, and strong shadows can still result in incorrect classifications. Furthermore, the dataset's bias towards adult dentition limits generalisability to juvenile and geriatric populations, and the dependence on 2D imaging restricts depth measurement.

The future work includes by adding multimodal dental imaging, growing clinically verified datasets, enabling longitudinal lesion tracking, and bolstering explainable AI for preventive treatment planning, future research will concentrate on enhancing ToothAid's clinical dependability and deployment. Large scale real world usage will be supported by additional optimisation for ultra low power edge devices and incorporation with tele-dentistry platforms while adhering to regulatory and healthcare privacy norms.

REFERENCES

[1] S. Patil, S. Albogami, J. Hosmani, S. Mujoo, M. A. Kamil, M. A. Mansour, H. N. Abdul, S. Bhandi, and S. S. S. J. Ahmed, "Artificial Intelligence in the Diagnosis of Oral Diseases: Applications and Pitfalls," *Diagnostics*, vol. 12, no. 5, p. 1029, 2022.

[2] N. Ensmenger and N. J. Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Achievements*, Cambridge University Press, 2010.

[3] I. H. Sarker, "Deep Learning: A Comprehensive Overview on Techniques, Taxonomy, Applications, and Research Directions," *SN Computer Science*, vol. 2, p. 420, 2021.

[4] J.-H. Lee, D.-H. Kim, S.-N. Jeong, and S.-H. Choi, "Detection and diagnosis of dental caries using a deep learning-based convolutional neural network algorithm," *Journal of Dentistry*, vol. 77, pp. 106–111, 2018.

[5] F. Casalegno, T. Newton, R. Daher, M. Abdelaziz, A. Lodi-Rizzini, F. Schürmann, I. Krejci, and H. Markram, "Caries detection with near-infrared transillumination using deep learning," *Journal of Dental Research*, vol. 98, no. 11, pp. 1227–1233, 2019.

[6] A. G. Cantu, S. Gehrung, J. Krois, A. Chaurasia, J. G. Rossi, R. Gaudin, K. Elhennawy, and F. Schwendicke, "Detecting caries lesions of different radiographic extension on bitewings using deep learning," *Journal of Dentistry*, vol. 100, p. 103425, 2020.

[7] H. J. Chang, S. J. Lee, T. H. Yong, N. Y. Shin, B. G. Jang, J. E. Kim, K. H. Huh, S. S. Lee, M. S. Heo, S. C. Choi, et al., "Deep learning hybrid method to automatically diagnose periodontal bone loss and stage periodontitis," *Scientific Reports*, vol. 10, no. 1, p. 7531, 2020.

[8] R. P. Danks, S. Bano, A. Orishko, H. J. Tan, F. Moreno Sancho, F. D'Aiuto, and D. Stoyanov, "Automating periodontal bone loss measurement via dental landmark localisation," *International Journal of Computer Assisted Radiology and Surgery*, vol. 16, pp. 1189–1199, 2021.

[9] F. O. Ozden, O. Ozgonenel, B. Ozden, and A. Aydogdu, "Diagnosis of periodontal diseases using different classification algorithms: A preliminary study," *Nigerian Journal of Clinical Practice*, vol. 18, p. 416, 2015.

[10] Y. Nakano, N. Suzuki, and F. Kuwata, "Predicting oral malodour based on the microbiota in saliva samples using a deep learning approach," *BMC Oral Health*, vol. 18, p. 128, 2018.

[11] 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.

[12] 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.

[13] J. A. Rodrigues, J. Krois, and F. Schwendicke, "Demystifying Artificial Intelligence and Deep Learning in Dentistry," *Brazilian Oral Research*, vol. 35, pp. 1–7, 2021.

[14] T. Hulsen, "Sharing Is Caring-Data Sharing Initiatives in Healthcare," *International Journal of Environmental Research and Public Health*, vol. 17, p. 3046, 2020.

[15] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 779–788.

[16] M. Abadi et al., "TensorFlow: A system for large-scale machine learning," in *12th USENIX symposium on operating systems design and implementation (OSDI 16)*, 2016, pp. 265–283.