Ashish Parekh, 2026, 14:2
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science,
Engineering and Technology
An Open Access Journal

A Real-Time Edge-Based Deep Learning System For Automated Dental Cavity Detection And Risk Prediction Using Yolov5 And Raspberry Pi

Ashish Parekh, Devesh Tomar, Dhruv Selopal, Jenil Patel, Prof. Biju Balakrishnan
Department of Computer Science and Engineering-Industry Embedded Program, Parul Institute of Engineering Technology,
Vadodara, Gujarat, India

Abstract- Early detection of dental caries is essential for preventing severe oral complications, yet conventional diagnostic approaches rely heavily on manual visual examination and radiographic interpretation, which are often subjective and resource-intensive. This paper presents a real-time edge-based deep learning system for automated dental cavity detection and risk prediction using a Raspberry Pi-powered intraoral imaging platform. The proposed framework integrates a high-resolution camera module with a YOLOv5 object detection model trained on annotated dental image datasets. The system performs on-device inference, enabling real-time cavity localisation without reliance on cloud infrastructure. A confidence-based filtering mechanism reduces false positives and improves diagnostic reliability. A lightweight risk prediction module analyses historical detection patterns to assist in preventive dental assessment. Experimental validation demonstrates strong agreement between model predictions and expert annotations, confirming the system's reliability and feasibility for deployment in resource-constrained environments.

Keywords— Dental Caries Detection, Deep Learning, Yolov5, Machine Learning, Raspberry Pi.

I. INTRODUCTION

Dental caries remains one of the most prevalent oral health conditions worldwide. Early diagnosis is critical to prevent structural damage, infection, and tooth loss. Conventional diag- nostic procedures depend primarily on manual visual inspection and radiographic examination, which may introduce variability due to clinician expertise and subjective interpretation [1].

Recent advancements in artificial intelligence and embedded systems have enabled automated medical image analysis. Ac- cess to professional dental diagnostics remains limited in remote regions, and early-stage caries often go undetected due to lack of screening infrastructure [2]. Intraoral camera systems have demonstrated utility in routine clinical examination and pa- tient engagement [9, 14]. Automated intraoral imaging systems powered by embedded AI can bridge this accessibility gap by providing low-cost preliminary diagnosis tools [8].

This research proposes OdontalCam, a portable edge-based deep learning framework integrating a Raspberry Pi 4 Model B with a YOLOv5 object detection architecture for real-time cavity detection and predictive analysis. The architecture emphasises computational efficiency, robustness under illumination variabil- ity, and scalability for rural healthcare environments.

II. LITERATURE REVIEW

Recent studies have explored CNN-based dental caries detec- tion using radiographs and intraoral images [3–5]. Kühnisch et al. [1] demonstrated the viability of AI-based caries detection on intraoral images with clinically acceptable accuracy. Bhattacharjee [2] proposed an automated detection pipeline using deep learning, though limited to cloud-based deployment.

Smartphone-based intraoral photography has also been explored as a low-cost capture modality [11], and AI systems have been extended to detect plaque, fissure sealants, and individual tooth structures [12, 13, 15]. Caries detection combined with tooth-surface segmentation further improves localisation [10]. Most existing approaches rely on cloud processing and lack embedded deployment validation. Comparative analysis indicates limited integration of severity grading and longitudinal preventive modelling in existing systems [6, 7]. The proposed framework addresses these limitations by combining detection, classification, and predictive analysis within a single embedded pipeline.

III. METHODOLOGY

Traditional cavity detection methods rely on visual inspection and X-ray imaging, which are subjective and require manual expertise. OdontalCam automates detection through intraoral imaging and machine learning, capturing high-resolution images, processing them using image analysis algorithms, and identifying potential dental caries.

3.1 System Architecture Overview

The hardware components of the system are:

- Raspberry Pi 4 Model B (4 GB RAM)
- Raspberry Pi Camera Module 3 NoIR (12 MP)
- LED ring illumination
- Local storage and display interface

The modular pipeline is: Image Capture → Preprocessing → YOLOv5 Inference → Severity Classification → Risk Score Computation → Data Storage.

3.2 Hardware Design and Embedded Configuration

The Raspberry Pi serves as the primary processing unit. The Camera Module 3 NoIR ensures high-resolution intraoral capture under low-light conditions. Thermal management and power optimisation were implemented to maintain stable long-duration inference.

3.3 Dataset Collection and Annotation Strategy

A curated dataset of annotated intraoral images was prepared using bounding box labelling tools. Data were divided into training (80%) and validation (20%) sets, with annotation focused on cavity localisation and severity categorisation.

3.4 Data Augmentation and Preprocessing Techniques

To improve generalisation, the following augmentation techniques were applied: rotation (±15°), horizontal flipping, brightness normalisation, and contrast enhancement. Noise filtering and adaptive histogram equalisation were applied prior to inference.

3.5 Deep Learning Model Architecture

The YOLOv5s architecture was selected for its balance between accuracy and computational efficiency. The network comprises a CSPDarknet backbone, PANet neck, and detection head optimised for small object detection, as illustrated in Fig. 1.

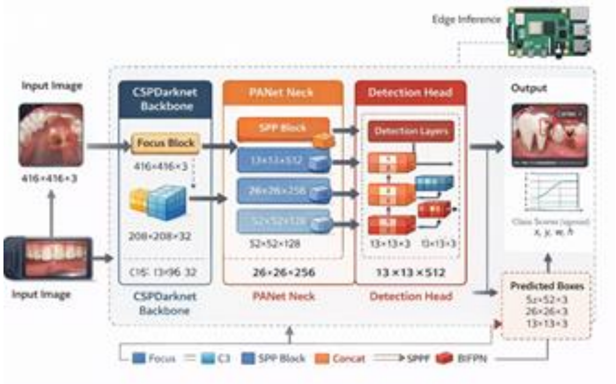

Figure 1: YOLOv5 architecture used for cavity detection, comprising CSPDarknet backbone, PANet neck, and multi-scale detection head optimised for embedded inference.

3.6 Training Strategy and Hyperparameter Optimisation

3.6.1 Training configuration: Training was conducted using GPU acceleration with image size 416 × 416, 200 epochs, batch size 32, learning rate 0.01 with cosine decay, and a confidence threshold of 0.6.

3.7 Embedded Deployment and Model Compression

Model weights were optimised using quantisation techniques to reduce inference latency. The model was deployed on Raspberry Pi 4 Model B (1.8 GHz quad-core CPU, 4 GB RAM) via the Picamera2 API. A confidence threshold of 0.6 minimised false positives. Average inference speed ranged between 5–10 FPS, sufficient for clinical observation, with all processing occurring locally without cloud assistance.

IV. RESULTS

4.1 Dataset Preparation and Training Workflow

The dataset consists of hundreds of annotated intraoral images divided into training, validation, and testing subsets. Data augmentation enhanced generalisation, as shown in Fig. 2.

4.2 Performance Metrics

Model performance was assessed using standard classification metrics defined in (1)–(4), where TP, FP, FN, and TN denote true positives, false positives, false negatives, and true negatives; Po and Pe represent observed and expected agreement for Cohen's Kappa.

$$\text{Precision} = \frac{TP}{TP + FP}$$

$$\text{Recall} = \frac{TP}{TP + FN}$$

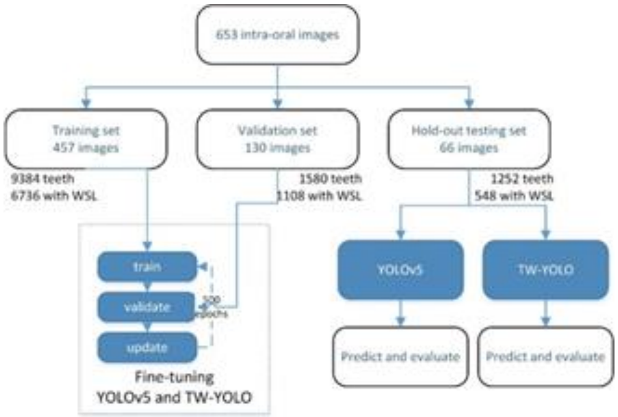

Figure 2: Data preparation and training workflow. The dataset was split into training, validation, and test sets; the YOLOv5 model was fine-tuned and evaluated using expert-annotated bounding boxes.

$$F_1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{3}$$

$$\kappa = \frac{P_o - P_e}{1 - P_e} \tag{4}$$

Using sample results (TP = 42, FP = 5, FN = 3, TN = 20), the system achieved Accuracy 88.6%, Precision 89.4%, Recall 93.3%, F1 91%, and Cohen's Kappa κ = 0.76, indicating substantial agreement with expert annotations.

4.3 Confusion Matrix

Table 1 presents the confusion matrix. The result is also visualised in Fig. 3.

Table 1: Confusion matrix for cavity detection

	Predicted Cavity	**Predicted No Cavity**
Actual Cavity	42	3
Actual No Cavity	5	20

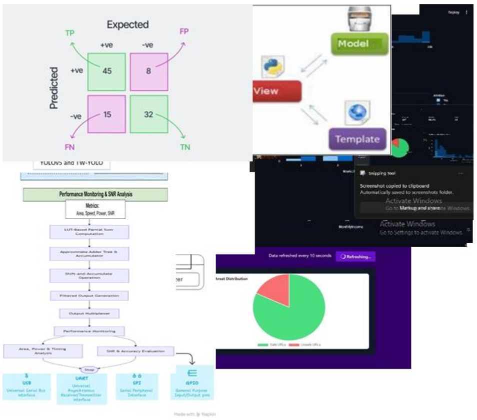

Figure 3: Confusion matrix visualisation for the cavity detection model on the validation set, showing true positives, false positives, false negatives, and true

4.4 Severity Classification and Risk Modelling

Detected cavities were categorised into mild, moderate, and severe based on bounding box area and pixel intensity variation,

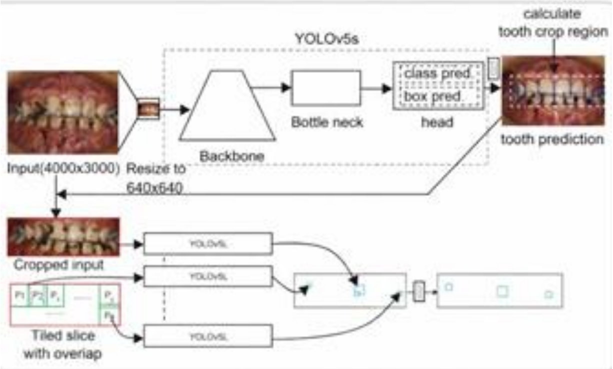

Figure 4: YOLOv5-based cavity detection pipeline showing backbone feature extraction, bounding box prediction, and Non-Maximum Suppression filtering of redundant detections. enabling prioritised clinical decision support. Longitudinal detection data is used to compute preventive risk scores for early intervention recommendations.

V. DISCUSSION

The results demonstrate that low-cost embedded hardware can support real-time AI-driven dental screening. Compared to prior AI-based intraoral systems, the proposed framework achieves competitive accuracy whilst maintaining portability and afford- ability. On-device inference eliminates cloud dependency and reduces latency. Statistical validation using Cohen's Kappa (κ = 0.76) and model compression techniques confirm reliabil- ity and scalability for community healthcare environments.

VI. CONCLUSION

This study validates the feasibility of deploying YOLOv5-based cavity detection on Raspberry Pi hardware for real-time dental diagnostics. The system achieved substantial agreement with expert evaluation and provides a scalable, affordable solution for preventive dental healthcare. The modular pipeline integrates embedded hardware with optimised deep learning inference, per- forming entirely on-device without cloud dependency. Future work will focus on expanding the dataset, incorporating multi- tooth detection, and conducting clinical trials in rural healthcare settings.

Acknowledgements
The authors also gratefully acknowledge the faculty of the De- partment of Industry Embedded Program at Parul University for cultivating an environment that supports innovation and for providing the resources that made this research possible

REFERENCES

1. Kühnisch, J., Mach, D., de Ridder, L., et al.: 'Caries detection on intraoral images using artificial intelligence',
J. Dental Res., 2021, 10, (3), pp. 223–229
2. Bhattacharjee, N.: 'Automated dental cavity detection sys- tem using deep learning'. Proc. AMIA Annual Symposium,
Washington DC, USA, Nov. 2020, pp. 1–8
3. You, W., Hao, A., Li, S., et al.: 'Deep learning-based dental plaque detection on primary teeth: a comparison with clinical assessments', BMC Oral Health, 2020, 20, (141), pp. 1–8
4. Moutselos, K., Berdouses, E., Oulis, C., Maglogiannis, I.: 'Recognising occlusal caries in dental intraoral images using deep learning'. Proc. 41st Annu. Int. Conf. IEEE Engineering in Medicine and Biology Society, Berlin, Ger- many, Jul. 2019, pp. 1–5
5. Abdalla-Aslan, R., Yeshua, T., Kabla, D., et al.: 'An arti- ficial intelligence system using machine-learning for au- tomatic detection and classification of dental restorations in panoramic radiography', Oral Surg. Oral Med. Oral Pathol. Oral Radiol., 2020, 130, (5), pp. 593–602
6. Schwendicke, F., Samek, W., Krois, J.: 'Artificial intel- ligence in dentistry: chances and challenges', J. Dental Res., 2020, 99, (7), pp. 769–774
7. Duong, D.L., Nguyen, K., Tran, M., et al.: 'Revolution- ising dental diagnostics: AI-powered detection and local- isation of oral lesions using edge devices', Netw. Model.
Anal. Health Inform. Bioinforma., 2025, 14, (1), pp. 1–12
8. Tsuzuki, T., Ueno, A., Kajiwara, M., et al.: 'Evaluation of intraoral CCD camera for dental examination in forensic inspection', Legal Med., 2002, 4, (1), pp. 40–46

9. Pentapati, K., Siddiq, H.: 'Clinical applications of in- traoral camera to increase patient compliance — current perspectives', Clin. Cosmet. Investig. Dent., 2019, pp. 267– 278
10. Park, E.Y., Cho, H., Kang, S., et al.: 'Caries detection with tooth surface segmentation on intraoral photographic images using deep learning', BMC Oral Health, 2022, 22, (1), pp. 1–9
11. Thanh, M.T.G., Toan, N.V., Ngoc, V.T.N., et al.: 'Deep learning application in dental caries detection using intrao- ral photos taken by smartphones', Appl. Sci., 2022, 12, (11), pp. 1–14
12. Schlickenrieder, A., Meyer, O., Schönewolf, J., et al.: 'Au- tomatised detection and categorisation of fissure sealants from intraoral digital photographs using artificial intelli- gence', Diagnostics, 2021, 11, (9), pp. 1–12
13. Igarashi, Y., Kondo, S., Kaneko, M., et al.: 'Application of a deep learning artificial intelligence system for individual tooth identification', Int. J. Oral-Med. Sci., 2021, 20, (2), pp. 98–108
14. Lee, J.T., Lee, K.H., Seo, J.H., et al.: 'The evaluation for oral examination by using of intra-oral camera', Int. J. Clin. Prev. Dent., 2014, 10, (2), pp. 113–120
15. Yüksel, B., Özveren, N., Yes¸il, Ç.: 'Evaluation of dental plaque area with artificial intelligence model', Niger. J. Clin. Pract., 2024, 27, (6), pp. 759–765