

Deep Learning Based Dental Cavity Detection

1 Akshaya.N, 2 Pauravi.B, 3 Dharsani.TS, 4 Dharaneesh.A, 5 Vidhyabharathi.T
akshayanatarajan2911@gmail.com, dharshusankar7@gmail.com

Department of Artificial Intelligence and Data Science, KPR College of Arts Science and Research, India

ABSTRACT

Abstract—Dental cavities are one of the most common oral health problems, and early identification can support timely dental care. This study presents a deep learning based dental cavity detection system using the Ultralytics YOLO11n (YOLOv11 Nano) object detection model. A dental image dataset containing annotated cavity and normal samples was prepared and divided into training, validation, and test sets. The model was trained in Google Colab with image preprocessing and augmentation to improve its ability to detect cavity regions. Model performance was evaluated using precision, recall, mAP50, and mAP50-95. The trained model was then used to predict cavity regions in unseen dental images. The proposed approach provides an automated and efficient method for preliminary cavity detection and can assist dental professionals as a decision-support tool. Future work can focus on larger and more diverse datasets, improved annotation quality, and real-world clinical validation. **Keywords**—dental cavity detection; deep learning; YOLO11n; object detection; dental images; computer vision

INTRODUCTION

Dental caries, commonly known as dental cavity, is one of the most common oral health problems affecting people of different age groups. If left untreated, it can cause pain, infection, tooth damage, and eventually tooth loss. Therefore, early and accurate detection of dental cavities is important for effective treatment and prevention of further complications. Traditionally, dental cavities are identified through clinical examination and dental imaging by dental professionals. However, manual examination can be time-consuming, and early-stage cavities may be difficult to identify accurately. Recent advancements in Artificial Intelligence (AI) and Deep Learning have provided new opportunities for automated analysis of dental images. In particular, object detection models can identify abnormal regions and determine their locations within an image.

This research proposes an **AI-Powered Dental Cavity Detection and Localization Using YOLO** system. The proposed approach uses a labeled dental image dataset to train a YOLO-based deep learning model for detecting and localizing cavity regions. The trained model can analyze dental images and identify suspected cavities using bounding boxes and confidence scores.

The proposed system aims to support faster and more consistent preliminary dental screening while reducing the dependency on completely manual image inspection. The performance of the model can be evaluated using metrics such

as precision, recall, and mean Average Precision (mAP). These metrics help determine the effectiveness of the model in accurately detecting and localizing dental cavities.

In addition, an automated dental cavity detection system can be useful for preliminary screening and early identification of suspicious tooth regions. By providing quick and consistent predictions, the system can help dental professionals focus their attention on areas that require further examination. This approach can also support the development of intelligent and accessible dental healthcare solutions using modern AI technologies.

METHODOLOGY


Dataset Preparation

The dataset used in this study contains **418 dental images** belonging to two classes: **cavity** and **normal**. The dataset was divided into **287 images for training**, **93 images for validation**, and **38 images for testing**. The **training set (287 images)** was used to train the YOLO11n model, while the **validation set (93 images)** was used to monitor and optimize the model during training. The **testing set (38 images)** was used to evaluate the model's performance on unseen dental images.

The available annotations were converted into **YOLO-compatible bounding-box format**, containing the class ID and normalized bounding-box coordinates. The class labels were assigned as **0 for cavity** and **1 for normal**. The annotation files were organized according to their corresponding image folders to ensure proper image-label association.

After dataset preparation, a YAML configuration file was created containing the paths of the training, validation, and testing datasets along with the class information. The prepared dataset was then used to train the **YOLO11n model for 50 epochs**. The trained model was evaluated using **Precision, Recall, mAP@50, and mAP@50-95** metrics.

YOLO11N Model



The proposed system uses a YOLO-based architecture for detecting and localizing dental cavities from input images. The architecture is divided into three major components: **Backbone**, **Neck**, and **Head**. The input image is initially resized to **640 × 640 pixels** and passed through the Backbone for feature extraction.

The **Backbone** consists of convolutional layers and C3K2 blocks that progressively extract important features such as edges, shapes, textures, and cavity patterns from the dental images. As the network becomes deeper, the spatial dimensions of the feature maps are reduced while the number of extracted features increases. The **SPPF (Spatial Pyramid Pooling Fast)** and **C2PSA** modules further enhance the extracted features by capturing information at different scales and improving feature representation.

The **Neck** is responsible for combining features from different levels of the Backbone. It uses **upsampling and concatenation operations** to merge low-level and high-level features. This multi-scale feature fusion helps the model detect cavities of different sizes and improves the localization of small cavity regions.

The **Head** forms the final detection stage of the architecture. It receives the refined features from the Neck and generates predictions for the detected regions. In the given architecture, multiple detection levels are used to perform multi-scale prediction, allowing the model to identify cavity regions at different resolutions. The final output provides the detected region along with its corresponding prediction information.

Overall, the combination of **Backbone for feature extraction**, **Neck for multi-scale feature fusion**, and **Head for final prediction** enables the proposed YOLO model to efficiently detect and localize dental cavities with good accuracy and speed.

Maintaining the Integrity of the Specifications

TRAINING AND PREDICTION

The model was trained in Google Colab using the Ultralytics framework. Training images and labels were supplied through the YOLO dataset configuration. After training, the best model weights were used to perform predictions on unseen test images. The predicted bounding boxes and confidence scores were used to identify cavity regions

Evaluation Metrics

1. Precision

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

Precision measures how many of the regions predicted as dental cavities are actually correct. A higher precision value indicates fewer false-positive cavity detections.

2. Recall

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

Recall measures how effectively the model detects the actual cavity regions present in the images. A higher recall indicates that fewer real cavity regions are missed.

3. mAP@50

$$mAP@50 = \frac{1}{N} \sum_{i=1}^N A P_i@0.50$$

mAP@50 represents the mean Average Precision at an **IoU threshold of 0.50**. It evaluates how accurately the model detects and localizes cavity regions when the predicted and actual bounding boxes have sufficient overlap.

4. mAP@50-95

$$mAP@50 - 95 = \frac{1}{10} \sum_{t=0.50}^{0.95} A P@t$$

mAP@50-95 calculates the average precision across **IoU thresholds from 0.50 to 0.95**, with a step of 0.05. It provides a stricter and more comprehensive measure of the model's detection and bounding-box localization performance.

Where: TP = True Positive, FP = False Positive, FN = False Negative, AP = Average Precision, and IoU = Intersection over Union.

RESULTS AND DISCUSSION


The proposed dental cavity detection system was implemented using the **YOLO11n** object detection model and trained for **50 epochs**. The performance of the trained model was evaluated

using Precision, Recall, mAP@50, and mAP@50–95 metrics. The model achieved a **Precision of 72.8%**, indicating that most of the regions predicted as cavities were correctly identified. The **Recall of 76.3%** shows that the model was able to detect a considerable proportion of the actual cavity regions in the test images. The obtained **mAP@50 of 78.7%** demonstrates good detection performance at an IoU threshold of 0.50. The **mAP@50–95 of 42.0%** indicates the model's performance across a range of IoU thresholds from 0.50 to 0.95. Overall, the results demonstrate that the proposed YOLO11-based approach can effectively detect dental cavity regions from dental images.

I. Table: Performance Evaluation of YOLO11n

Metric	Value
Model	YOLO11n
Training Epochs	50
Precision	72.8%
Recall	76.3%
mAP@50	78.7%
mAP@50–95	42.0%

BAR CHART



.LIMITATIONS

Although the proposed YOLO11n-based dental cavity detection system achieved promising results, it has certain limitations. The performance of the model depends on the quality, quantity, and diversity of the dental image dataset used for training. Variations in image quality, lighting conditions, tooth orientation, and image acquisition devices may affect detection performance. The model may also produce false positive or false negative detections in complex dental images. Furthermore, the current system focuses mainly on detecting cavity regions and does not provide detailed information about the severity or stage of dental caries. The achieved mAP@50–95 of 42.0% also indicates that the accuracy of precise bounding-box localization can be further improved. The model may require further optimization to improve its performance on unseen dental images from different sources. Real-time deployment may also be affected by computational and hardware requirements. The current system is designed primarily for cavity detection and may not cover other dental conditions. In addition, variations in dental image acquisition environments may affect the generalization of the model. Future improvements can focus on using larger and more diverse

datasets, improving annotation quality, optimizing the model architecture, and integrating the system with clinical

CONCLUSION

The proposed dental cavity detection system successfully demonstrates the effectiveness of the **YOLO11n object detection model** for automated identification of dental cavity regions from dental images. The model was trained for **50 epochs** and achieved a **Precision of 72.8%, Recall of 76.3%, mAP@50 of 78.7%, and mAP@50–95 of 42.0%**. The obtained **mAP@50 of 78.7%** indicates that the proposed model provides good cavity detection performance at an IoU threshold of 0.50, while the Recall value of **76.3%** shows that the model is capable of identifying a considerable proportion of the actual cavity regions. The system also successfully performed prediction on new dental images and provided bounding-box-based cavity detection with confidence scores.

The proposed approach can reduce the effort required for preliminary examination of dental images and may assist dental professionals in identifying suspicious cavity regions more efficiently. The lightweight **YOLO11n** architecture also provides a suitable foundation for developing an automated and practical dental image analysis system. However, further improvement is required to enhance precise localization, particularly considering the **mAP@50–95 value of 42.0%**. Future work can focus on increasing the size and diversity of the dataset, improving annotation quality, optimizing model parameters, and incorporating more advanced image-processing techniques. The system can also be extended to identify different stages of dental caries and other dental conditions. With further validation using larger and clinically diverse datasets, the proposed system has potential for integration into computer-aided dental screening and decision-support applications.

ACKNOWLEDGEMENT

We sincerely express our gratitude to the management, principal, faculty members, and Department of Artificial Intelligence and Data Science of KPR College of Arts Science and Research for providing the academic environment and support required to complete this project. We thank our project guide and teaching faculty for their valuable suggestions, technical guidance, encouragement, and continuous support throughout dataset preparation, model development, training, and evaluation. We also acknowledge the resources and open-source tools that made the implementation possible, including Google Colab and the Ultralytics YOLO framework. Finally, we thank everyone who directly or indirectly supported the successful completion of this research work.

REFERENCES

- [1] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 779–788, 2016.
- [2] J. Redmon and A. Farhadi, "YOLO9000: Better, Faster, Stronger," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7263–7271, 2017.

- [3] J. Redmon and A. Farhadi, "YOLOv3: An Incremental Improvement," *arXiv preprint arXiv:1804.02767*, 2018.
- [4] A. Bochkovskiy, C. Y. Wang, and H. Y. M. Liao, "YOLOv4: Optimal Speed and Accuracy of Object Detection," *arXiv preprint arXiv:2004.10934*, 2020.
- [5] C. Y. Wang, A. Bochkovskiy, and H. Y. M. Liao, "YOLOv7: Trainable Bag-of-Freebies Sets New State-of-the-Art for Real-Time Object Detectors," *arXiv preprint arXiv:2207.02696*, 2022.
- [6] A. Wang, H. Chen, L. Liu, K. Chen, Z. Lin, J. Han, and G. Ding, "YOLOv10: Real-Time End-to-End Object Detection," *arXiv preprint arXiv:2405.14458*, 2024.
- [7] G. Jocher and J. Qiu, "Ultralytics YOLO11," Ultralytics, version 11.0.0, 2024.
- [8] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks," *Advances in Neural Information Processing Systems*, vol. 28, 2015.
- [9] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 6, pp. 1137–1149, 2017.
- [10] W. Liu et al., "SSD: Single Shot MultiBox Detector," *European Conference on Computer Vision (ECCV)*, 2016.
- [11] M. Schwendicke, J. Samek, and G. Krois, "Artificial Intelligence in Dentistry: Chances and Challenges," *Journal of Dental Research*, vol. 99, no. 7, pp. 769–774, 2020.
- [12] M. Schwendicke, G. Krois, and J. Samek, "Artificial Intelligence in Dentistry: A Systematic Review," *Journal of Dental Research*, 2020.
- [13] S. M. Khanagar et al., "Developments, Application, and Performance of Artificial Intelligence in Dentistry—A Systematic Review," *Journal of Dental Research*, 2021.
- [14] M. C. Schwendicke et al., "Deep Learning for Caries Detection: A Systematic Review," *Journal of Dentistry*, vol. 122, p. 104115, 2022.
- [15] A. Moharrami et al., "Detecting Dental Caries on Oral Photographs Using Artificial Intelligence: A Systematic Review," *Oral Diseases*, vol. 30, no. 4, pp. 1765–1783, 2024.
- [16] S. Negi et al., "Artificial Intelligence in Dental Caries Diagnosis and Detection: An Umbrella Review," *Clinical and Experimental Dental Research*, vol. 10, no. 4, e70004, 2024.
- [17] A. N. Uddin et al., "Applications of AI-Based Deep Learning Models for Detecting Dental Caries on Intraoral Images: A Systematic Review," *Evidence-Based Dentistry*, vol. 26, pp. 71–72, 2025.
- [18] L. P. Abbott, A. Saikia, and R. P. Anthonappa, "Artificial Intelligence Platforms in Dental Caries Detection: A Systematic Review and Meta-Analysis," *Journal of Evidence-Based Dental Practice*, vol. 25, no. 1, p. 102077, 2025.
- [19] Q. T. Lam et al., "Evaluating YOLO for Dental Caries Diagnosis: A Systematic Review and Meta-Analysis," *Evidence-Based Dentistry*, vol. 26, p. 176, 2025.
- [20] P. Mongkolwat et al., "Clinical Application of Deep Learning for Enhanced Multistage Caries Detection in Panoramic Radiographs," *Scientific Reports*, vol. 15, article 33491, 2025.
- [21] "Detection of Dental Caries Under Fixed Dental Prostheses by Analyzing Digital Panoramic Radiographs with Artificial Intelligence Algorithms Based on Deep Learning Methods," 2025.
- [22] D. Ozdemir et al., "An Enhanced Deep Learning Model for Detection and Classification of Dental Caries in Panoramic Radiographs," *Neural Computing and Applications*, vol. 38, article 3, 2026.
- [23] "Deep Learning-Based Dental Caries Diagnosis on Panoramic Radiographies: Performance of YOLOv8 Versus Human Observers," *Diagnostics*, vol. 16, no. 8, p. 1150, 2026.
- [24] A. N. Uddin et al., "Applications of AI-Based Deep Learning Models for Detecting Dental Caries on Intraoral Images: A Systematic Review," *Evidence-Based Dentistry*, 2025.
- [25] "Artificial Intelligence for Dental Caries Detection: An Umbrella Review," *Journal of Dentistry*, vol. 175, p. 106944, 2026. .

