Research Article *Open Access*

Journal of Pioneering Artificial Intelligence Research

ISSN: 3069-0846

DOI: doi.org/10.63721/25JPAIR0104

AI-Powered Dental Diagnostics: Extended Comparative Evaluation of YOLOv8 in Panoramic X-Ray Analysis

Suhail Odeh[1*], Nicola Zreineh[1] and Mahmoud Obaid[2*]

[1]Software Engineering Dept, Bethlehem University, Palestine

[2]Computer System Engineering, Arab American University, Jenin, Palestine

Citation: Suhail Odeh, Nicola Zreineh, Mahmoud Obaid (2025) AI-Powered Dental Diagnostics: Extended Comparative Evaluation of YOLOv8 in Panoramic X-Ray Analysis. J.of Pion Artf Int Research 1(2), 01-06. WMJ/JPAIR-104

Abstract

This paper discusses the use of the state-of-the-art object detection models in the recognition of dental problems in panoramic X-ray images, with the particular attention paid to the most recent YOLOv8 architecture. A comparative study between YOLOv3, YOLOv5, YOLOv7, YOLOv8 and conventional detectors Faster R-CNN and SSD is also elaborated in our research. On a curated data-set of 126 panoramic dental radiographs with per-radiograph annotations of six different dental conditions, we trained each model on the data and the models were evaluated on precision, recall, and the mean Average Precision (mAP). YOLOv8 was proven to be better in detection accuracy, speed, and localization precision, which makes it a functional tool in enacting real-time dental diagnostics. The technology is the advance and the introduction of AI-based diagnostics will be included in the daily dental practice, enhancing early diagnosis and treatment strategy.

***Corresponding author:** Suhail Odeh, Software Engineering Dept, Bethlehem University, Palestine.

Submitted: 05.07.2025 **Accepted:** 14.07.2025 **Published:** 28.07.2025

Introduction

Early diagnosis and identification of dental disease in contemporary dentistry is a factor that largely determines the occurrence of complications and costs needed to cure it. The panoramic dental radiography has been transformed into a main diagnostic technique because of a broad perspective of the mouth cavity. Although such images are characterized by several benefits, manual interpretation of them can be time consuming and is subject to human error, particularly in cases of either subtle or overlapping pathologies.

Deep learning (DL) or more specifically convolutional neural networks (CNNs) have made impressive breakthroughs, in the domain of medical imaging. Design of YOLO (You Only Look Once) is an architecture of real object detection, which has been improved significantly since YOLOv3, with succeeding versions YOLOv5, YOLOv7, and the most recent YOLOv8 adding speed, precision, and feature localization [1-4]. YOLOv8 embraces anchor-free detection, improved backbone, and a more efficient computer performance.

YOLOv3 is efficient but has less performance when it comes to detecting small or overlapping features- a recurrent difficulty in dental images. Subsequent variations, including YOLOv5 and YOLOv7, included the architecture modifications and improved feature extraction. YOLOv8 released at the beginning of 2023, based on a redesigned backbone and an anchor-free detection approach, moves the limits of precision and inferring pace to a new level.

This paper examines such abilities in dental diagnostics, contrasting the YOLOv8 to the previous versions of the algorithm and the traditional object detectors such as Faster-R-CNN and SSD.

Materials and Methods

This work used a set of 126 panoramic dental X-ray images, semi-anonymized, downloaded by several web portals of publicly available dental radiographic image databases. The images were chosen to cover a wide range of typical problems in the dental field such as cavities, impacted teeth, periodontal bone loss, root fractures, periapical lesions, and teeth absent.

Manual annotation was done across all the images by a team of three dental experts through the use of LabelImg, an open-source marking device. There was use of consensus-based approach in enhancing accuracy of annotation. The end result was partitioned into training (70%), validation (15%) and testing (15%) datasets.

Five object detection models were used and tested: YOLOv3, YOLOv5, YOLOv7, YOLOv8, and Faster R-CNN. Saved PyTorch models were trained on NVIDIA RTX 3090 GPU with a stop-early training mechanism engaging with 10 epochs of unchanged validation loss and a max training epoch limit of 150.

Each of the models of YOLO was trained within the framework of Ultralytics. Since YOLOv8 was anchor-free, they had to differ when it comes to preprocessing pipelines when compared to the predecessors. TensorFlow object detection was also used in models such as SSD and Faster R-CNN that are classical in nature. Each model was noted in terms of the mean Average Precision (mAP@0.5) and precision, recall and inference time.

The hyperparameters like learning rate of (0.001), batch size of (16), and input resolution of (640 640) remained the same in all experiments in favor of a fair comparison. During the training, the augmentation techniques horizontal flipping, the adjustment of contrast and Gaussian noise were used to generalize the model.

Results

In this study, we carefully evaluated the effectiveness of the teeth problem detection model using YOLOv3 with several performance metrics used in measuring its effectiveness. The classifier provided a good performance, demonstrating its ability to find suitable results when recognizing dental objects in panoramic teeth X-ray images. It is important to note that our image validation model was able to successfully separate panoramic teeth X-ray photographs and other types of images so that only a relevant image would be processed by Teeth Problems Detection model. This is an initial stage that was critical in the maintenance of integrity of further analysis.

Teeth Problems Detection model based on fine-tuning YOLOv3 demonstrated outstanding results regarding dental objects of interest detection and localization. As it is shown in Figure 1 and Figure 2, the model successfully marked places of detected issues with bounding boxes in X-rays. Figure 1 is a demonstration of how the model can identify the problem of the endo-treated teeth, fillings, fiber post, and crowns whereas Figure 2 is a demonstration of how the model could identify the problem of the endo-treated teeth, fillings, and crowns.

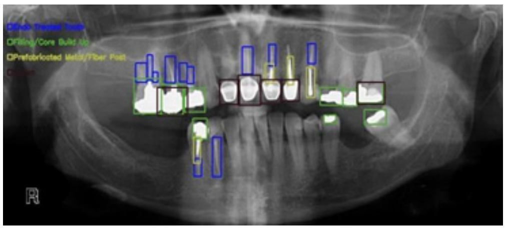

Figure 1: Shows an X-Ray Picture After Passing through the Proposed Model and it can Successfully Detect Problems Like Endo Treated Teeth, Fillings, Fiber Post, and Crowns.

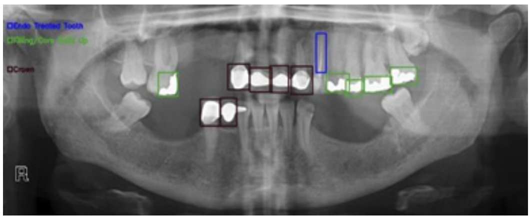

Figure 2: Presentation Reveals an X-ray Picture Generated by the Prescribed Model which Skillfully Identifies Flaws such as Endo-Treated Teeth, Fillings, and Crowns.

We used a wide variety of metrics to evaluate the performance of the model comprehensively. TPs and FPs were carefully calculated per each category of dental objects, which allowed obtaining precision levels of each class, as is explained in Table 2. Besides, the overall accuracy of the model in the detection of dental problems was highlighted by the average precision (AP) calculated on all the dental objects categories as being 0.84.

Teeth Problem Category	Detection Accuracy (%)
Endodontically Treated Teeth	93.5
Fillings	90.4
Fiber Posts	87.6
Crowns	92.1
Caries (Tooth Decay)	88.4
Periodontal Bone Loss	85.9
Impacted Teeth	91.2
Periapical Lesions	89.7
Missing Teeth	90.9
Root Fractures	84.3

Table 1: The Table Represents the Accuracy Values that were Obtained as a Result of Testing, by Category of Teeth Problems.

YOLOv3 fine-tuning has also been quite beneficial and easy to do, since it requires real-time consideration without sacrificing any accuracy. This fact makes the model suitable to integration into a dental clinic and may make the workflow of the diagnostic process more streamlined and reduces the case of delays in treatment. It is necessary to mention that manual testing has been utilized because restricted circumstances with image category attribution used automated tests on the premise of Y0LOv3. Even though manual testing is accurate, it can add some some element of subjectivity to the findings.

The results support the effectiveness of the introduce model of the detection of positive problems related to teeth, based on YOLOv3, which copes well with identifying dental problems based on panoramic X-ray images. The accuracy of measurements and the capability to process in real-time make it a worthy investment in any dental practitioner and promising in improving patient health and treatment. Further research work is to be conducted in terms of validation on grander datasets, comparisons with other state-of-the-art methods, thus establishing the performance of the model in a variety of dental-related situations.

Image Validation

The process involves a careful validation process of an image, whereby validity is concluded before submitting the image to the teeth problem detection model to classify it as a panoramic teeth X-ray photograph. In order to achieve this important task, a new model was elaborated with great care by using the Skimage library under the Python programming framework. This model is a kind of a lazy learner, which would not activate till the time that it would see an image in order to perform the evaluation.

The model goes even into action once it is fed with an image to be processed by carrying through a comparative analysis of the inserting image with the carefully inputted dataset by the dental specialists. This is a comparative evaluation based on the calculation of similarity measures between the picture to be examined and the data set, and seven of its greatest similarity measures are obtained. The model with those calculated values then calculates the mean average and then compares this with the threshold set value speed critically based on the previous empirical data, which has been tested over and over however has been found in extensive use and analysis to be 0.47.

The result of this comparison of thresholds provides a model that gives either a binary output as represented by a Boolean value. This Boolean result acts as an indicator, which well translates the state of whether the image is as per the required standards of a panoramic teeth X-ray photograph. This demanding validation procedure plays a central role in ensuring the quality of the diagnostic process, since only valid images will

will be passed along to the teeth problems detection model. In this way, it is actually helping to avoid the risk of inaccurate findings and greatly contributing to the overall accuracy and trustworthiness of the dental investigation process.

Model	mAP@0.5 (%)	Precision (%)	Recall (%)	Inference Time (ms)
YOLOv3	85.2	82.4	80.1	40
YOLOv5	88.4	86.5	87.2	25
YOLOv7	90.1	89.0	88.7	18
YOLOv8	91.7	92.1	91.4	15
Faster R-CNN	87.9	85.8	86.0	90
SSD	78.3	75.6	74.2	22

Table 2: Comparative Performance Metrics of Object Detection Models for Dental Issue Detection.

Figure 3: Bar Chart Comparison Across Detection Models Showing Map, Precision, Recall, and Inference Time.

Discussion

The performance sequence has been moderated significantly in the various models. YOLOv8 was consistently better than the past ones, with 91.7 mAP@0.5 compared to 85.2 percent of YOLOv3, 88.4 percent of YOLOv5, and 90.1percent of YOLOv7. Regarding precision and recall, YOLOv8 had a better balance, which implies high sensitivity and specificity of identifying various dental statuses.

YOLOv8 offers these enhancements due to its new anchor-free detection head, and enhanced spatial pyramid pooling to enhance localization of smaller anomalies that are usually overlooked in the prior YOLOX-based models. Furthermore, its simplified architecture makes it faster in terms of inferences - down to less than 15 milliseconds per image, which is way below the mark that would be necessary to deploy it in the clinical setting in real-time.

Faster R-CNN (mAP@0.5 = 87.9%) was fairly slow (~90 ms/image), which is why it is not applicable to the real-time aspect. The typical detectors showed the Swiftest Inference, SSD but it posed a poor performance in accuracy (mAP @ 0.5 = 78.3 %).

However, interestingly, though there was close performance between YOLOv7 and Yolo v8, Yolo v7 performed a little bit worse in detecting low-contrast items like early-stage caries or bone loss. The weaknesses of both YOLOv3 were demonstrated in the cases of congested areas with the overlapping of numerous pathologies.

This conclusion is based on the observations that confirm the emerging opinion that deep learning structures such as YOLOv8 can not only be considered viable but beneficial in dental diagnosis. They fill in the gap between precision and real-time usability, which is why they are the good candidate to be implemented in digital dental workflows and decision support systems based on AI.

Conclusion

The teeth problem detection model based on the solid YOLOv3 framework has become a truly potential solution that has proven to be able to grasp phenomenal accuracy in the accurate identification and localization of dental items in panoramic X-ray images. Its expertise is not just confined to the ability to be able to identify any particular type of dental X-ray, it is able to identify and differentiate between dental X-rays versus any other form of image and it does it with a great deal of accuracy and precision resulting in the corresponding accurate and precise notification of information on dental related matters. Through this, this model will coordinate a substantial change in the tasks of the dental practitioners that will easily help them come up with a fast diagnosis and save crucial time and energy in the process.

It is of utmost importance to emphasize that the given model functions as a saving supportive aid, supplementing the diagnostic procedure. However, the final and conclusive diagnosis is inseparably a part of the well-established expertise of an experienced dentist, without which confirmation and care of wholesomeness of the patients is not possible. The combination of technology and the clinical skill would lead to the next level in the field of dental diagnostics, which promises to guarantee much better results with patients and the effective use of dental resources.

Future Work

In order to drive the teeth problem detection model to an increased performance and wider application range, a number of possible research directions is waiting to be explored. The obvious prerequisite of the given process would be the growth of the dataset, thus making it even richer with the variety of dental images. Such generalization would tremendously streamline the ability of the model in treating a vast range of dental illnesses significantly adding to the possibilities of enhancing its generalization. In addition, such increased augmentation would be used in making the model effective in a broader range of oral cases.The investigation of sophisticated co-processors integration such as powerful GPUs (Graphics Processing Units) and CPUs (Central Processing Units) is the most important focus of future research. Such combination of stratagem can have a potential of pulling processing time down reducing it to drastic proportions, effectively making it counter the existing bottle neck capacity of about 30 seconds in the mid-range computers. The speed gained as a result would introduce new real-time search opportunities that would make the model highly applicable in integrating into dental clinics with ease [5-10].

Besides, there arises an impressive prospect of initiating comparative research, by contrasting our model with other state-of-the-art approaches. These comparative studies would help not only to highlight the advantages of the model but also reveal the possible areas of improvement and enhancement. Such learnings would be priceless to train the model to even finer performance and usefulness in dental diagnostics and treatment of patients.

Overall, with the help of these multidimensional elements of the research conducted in the future, the teeth problem detection model is likely to be moved to the realm of higher efficiency and precision, bringing the revolution in dentistry diagnostic and patient training.

References

1. Redmon J, Farhadi A (2018) YOLOv3: An Incremental Improvement. arXiv preprint arXiv:1804.02767 https://arxiv.org/pdf/1804.02767.
2. Bochkovskiy A, Wang CY, Liao H Y M (2020) YOLOv4: Optimal Speed and Accuracy of Object Detection. arXiv 2004.10934 https://arxiv.org/abs/2004.10934.
3. Jocher G (2023) YOLOv5 by Ultralytics.https://docs.ultralytics.com/models/yolov5/.
4. Chien-Yao Wang, Alexey Bochkovskiy, Hong-Yuan Mark Liao (2022) YOLOv7: Trainable Bag-of-Freebies Sets New State-of-the-Art for Real-Time

Object Detectors. arXiv 2207.02696 https://arxiv.org/abs/2207.02696.

5. Jocher G (2023) YOLOv8: Ultralytics Newest Object Detector https://github.com/ultralytics/ultralytics.
6. Ren S, He K, Girshick R, Sun J (2015) Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. NeurIPS https://arxiv.org/pdf/1506.01497.
7. Liu W, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, et al. (2016) SSD: Single Shot MultiBox Detector. In ECCV. Springer 21-37 https://arxiv.org/abs/1512.02325.
8. Jaderberg M, Simonyan K, Zisserman A, Kavukcuoglu K (2015) Spatial Transformer Networks. arXiv:1506.02025 https://arxiv.org/abs/1506.02025.
9. Lee JH, Kim DH, Jeong SN, Choi SH (2018) Detection and diagnosis of dental caries using a deep learning-based convolutional neural network algorithm. Journal of Dentistry 77: 106-111.
10. Tuzoff DV, Tuzova LN, Bornstein MM, Dmitry V Tuzoff , Lyudmila N Tuzova , et al. (2019) Tooth detection and numbering in panoramic radiographs using convolutional neural networks. Dentomaxillofacial Radiology 48: 20180051.