A Novel Artificial Intelligence Approach to Kennedy Classification for Partially Edentulous Patients Using Panoramic Radiographs

Keywords

Artificial Intelligence
YOLO
Panoramic Radiographs
Kennedy Classification
Partially Edentulous
YOLOv8

Authors

Nermeen Ahmed Hassan *
(BDs, MSc, PhD)

Alzahraa Abdelmongi §
(BSc)

Sara Magdi §
(BSc)

Mariam Shaltout §
(BSc)

Yara Aboelhasan §
(BSc)

Yasmine Elhariry §
(BSc)

Ensaf H. Mohamed ‡^
(BSc, MSc, PhD)

Address for Correspondence

Nermeen Ahmed Hassan *

Email:
nermeen.hassan@dentistry.cu.edu.eg

* *Lecturer, Department of Prosthodontics, Faculty of Dentistry, Cairo University, Cairo, Egypt.*

§ *Faculty of Computer and Artificial intelligence, Helwan University Cairo, Egypt.*

‡ *Associate Professor at School of Information Technology and Computer Science, Nile University, Giza, Egypt*

^ *Associate Professor at Department of Computer Science, Faculty of Computer and Artificial Intelligence, Helwan University, Cairo, Egypt*

Received: 13.09.2024
Accepted: 18.02.2025

doi: 10.1922/EJPRD_2801Hassan09

ABSTRACT

Objectives: This study aimed to develop an artificial intelligence system for automated classification of partially edentulous arches from panoramic radiographs using the Kennedy classification system and Applegate's rules, alongside identifying existing teeth for automated reporting. Methods: From 5261 anonymized digital panoramic radiographs collected from publicly available datasets, 1875 high-quality images were selected and divided into training (80%), validation (10%), and testing (10%) sets. Teeth were manually annotated on the Roboflow platform following the Universal Numbering System. To enhance model robustness, data augmentation techniques were applied, expanding the dataset to 2398 images. For tooth detection, a YOLOv8s deep learning model was trained for 80 epochs (batch size: 16, learning rate: 0.01). Performance was evaluated using precision, recall, F1 score, and mean average precision. Detected teeth were used to classify partially edentulous areas based on the Kennedy system. Modification areas were identified by analyzing detected and missing teeth, measuring bounded distances in millimetres, and classifying free-end saddle gaps. Results: The YOLOv8s model achieved a mean average precision (mAP50) of 98.1% for tooth identification, with precision and recall of 95.7% and 95.8%, respectively. For Kennedy classification, the model demonstrated precision of 0.962, recall of 0.931, and an F1-score of 0.939 across maxillary and mandibular arches. Conclusions: The high accuracy and efficiency of this AI-driven approach can standardize classification, reduce diagnostic variability, and alleviate the workload for dental professionals, enabling seamless integration into clinical practice. Clinical Relevance: This AI system provides a consistent, accurate, and reliable method for classifying partially edentulous arches from panoramic radiographs, reducing manual assessment variability, easing practitioner workload, and enabling large-scale analysis of partial edentulism prevalence.

INTRODUCTION

Partial edentulism presents diverse challenges in terms of functional, aesthetic, and psychological impacts on patients.[1] The classification of partially edentulous patients is an important component in the field of prosthodontics, guiding the treatment planning and rehabilitation process. Various classification systems exist for grouping several patterns of partial edentulism.[2,3] The Kennedy classification system, introduced by Dr. E. Kennedy,

remains the most widely recognized systems for categorizing partially edentulous arches. Its widespread adoption is attributed to its inherent simplicity and ease of use.[4,5] The Kennedy classification was then refined by Applegate's rules which are a set of guidelines that address specific scenarios, such as the determination of the primary edentulous area in complex cases and the inclusion or exclusion of third molars. This system effectively categorizes the vast array of partially edentulous arch configurations, estimated to number nearly 65,000, using a concise and unambiguous framework.[5]

The Kennedy classification system, when supplemented by Applegate's rules provides a robust foundation for treatment planning, communication among dental professionals, and the design of removable partial dentures (RPDs).[6] However, the application of Applegate's rules alongside the classification systems relies on subjective visual assessments by dental professionals, making them prone to variability in diagnostic accuracy and lacking the consistency needed for optimal patient management. Moreover, data on the global prevalence of partial edentulism classified by Kennedy classes remains scarce and potentially outdated. This scarcity likely stems from the significant workload placed on dental professionals during screening procedures, which often involve evaluating a large volume of panoramic radiographs.

In recent years, artificial intelligence (AI) has begun to revolutionize the field of dentistry. Machine learning algorithms are being utilized to analyse dental images with high accuracy. This technology has shown promising results when diagnosing dental caries and detecting the type of dental implant with results comparable or even superior to human dentists.[7,8] AI's potential extends beyond caries detection, with ongoing research exploring its application in diagnosing other oral conditions like periodontal disease and even oral cancer.[9,10,11] Additionally, AI is currently being implemented not only in traditional dental imaging but also in Cone Beam Computed Tomography (CBCT), further expanding its utility in advanced diagnostic and treatment planning processes.[12] This enhanced diagnostic capability of AI presents a significant opportunity to improve the classification of partially edentulous patients.

The aim of the current study was to develop and implement an AI system to automatically classify partially edentulous arches from panoramic radiographs based on Kennedy classification system and to identify and segment the existing teeth to provide automatic reporting.

MATERIAL AND METHODS

DATA PREPARATION AND PREPROCESSING

Data Collection

A total of 5261 anonymized digital panoramic radiographic images for partially edentulous and fully dentate arches were collected from publicly available datasets on Mendeley, Zenodo, and Figshare.[13,14,15]

Informed consent and Institutional Review Board (IRB) approval were not required due to the anonymized nature of the data, as outlined by the National Institutes of Health (NIH) Office of Human Subjects Research Protections (OHRP). Panoramic radiographs with significant noise, artifacts, blurring or those with deciduous teeth were excluded from the dataset. Therefore, 1875 anonymized digital panoramic radiographic images of high quality were included in the study as an initial dataset. Then, the dataset was divided into three distinct subsets: 80% training set, 10% validation set, and 10% testing set. The training set was comprised of 1500 images and serves as the foundation for model training.

Data Annotation

To facilitate the model's analysis of tooth presence and location, all teeth within each panoramic radiograph were meticulously manually labelled using the Roboflow annotation platform. This annotation process adhered to the established Universal Numbering System (UNS) for consistent and accurate tooth identification. The UN system begins by numbering tooth #1 for the upper right third molar. Numbering then proceeds in a clockwise direction around the dentition, attributing subsequent sequential numbers (#2, #3, etc.) to each tooth encountered until reaching tooth number #32, which corresponds to the lower right third molar.[16]

The annotation process involved outlining each tooth within the image using a bounding box, which is a characteristic feature of object detection. Object detection focuses on identifying and localizing objects within an image by enclosing them in bounding boxes, defined by precise coordinate pairs representing the upper-left and lower-right corners of the box. (*Figure 1*) In contrast, object segmentation involves delineating the exact shape and boundaries of objects at a pixel level, providing a more detailed representation. For this study, object detection was employed to efficiently localize and identify teeth, as it aligns with the model's goals and computational requirements.

Data Augmentation

To enhance the robustness and variability of the model, several data augmentation techniques were employed on the training dataset. These techniques included random rotations (±10°), random shear transformations (±10° horizontally and vertically), adjustments in contrast (between -10% and +10%), horizontal flipping of images, and the addition of minimal noise to simulate real-world variations in radiographs. By randomly applying these transformations, the number of the dataset was increased without introducing entirely new information. This process effectively expanded the training data to 2398 images.

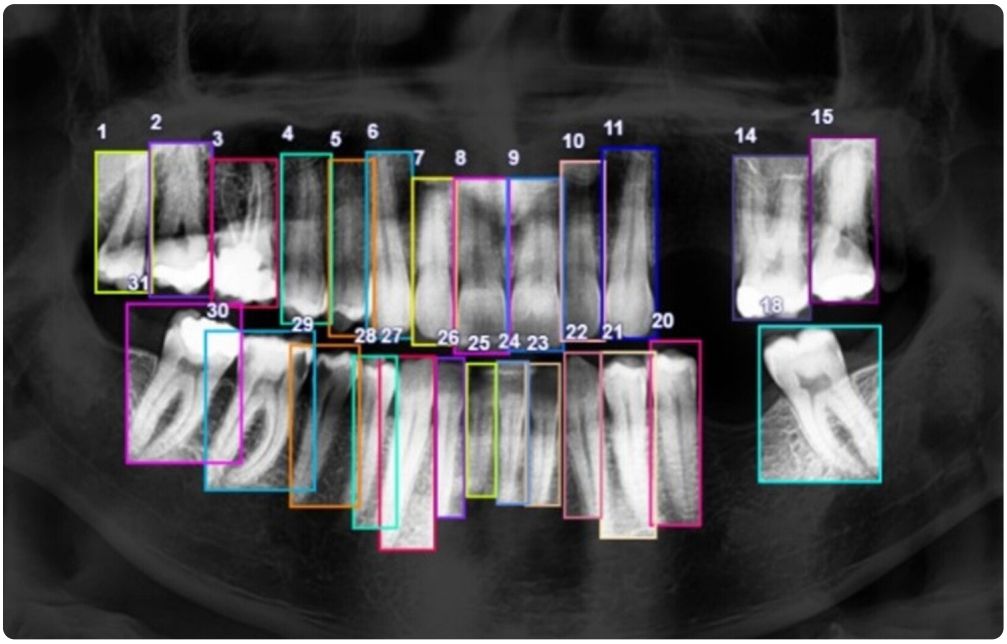

Figure 1: Annotation of the teeth on panoramic radiograph according to the universal numbering sytem.

TEETH DETECTION PHASE

Architecture Model

The proposed model for the current study employed a YOLOv8s deep learning architecture which is a single-stage convolutional neural network (CNN) object detection model.

YOLO, which stands for You Only Look Once, is a deep learning algorithm renowned for its speed and accuracy in object detection and segmentation tasks. As a single-stage algorithm, YOLO detects and segments objects in an image in a single pass, making it highly efficient and suitable for real-time applications. YOLO divides the image into a grid and predicts the presence of objects along with their bounding boxes within each grid cell.[17]

Model Training

The training process involved meticulously preparing the dataset for optimal utilization by the YOLOv8s model. All panoramic radiographs were first converted to a uniform JPEG format to ensure consistency in the data format fed into the model. Additionally, all images were resized to a fixed dimension of 640 x 640 pixels. This standardization ensures that the model receives images with consistent dimensions, simplifying the learning process. Moreover, a batch size of 16 was used and the model was trained for 80 epochs. The choice of batch size and number of epochs was determined based on preliminary experiments to balance computational efficiency and model performance. The learning rate was set at 0.01 and other hyperparameters were fine-tuned using the validation set.

Evaluation of Model Performance

During training, the model's performance was monitored using training and validation loss metrics. These metrics quantify the difference between the predicted and actual values, with lower losses indicating better performance. Training loss measures how well the model learns from the training data, while validation loss reflects its ability to generalize to unseen data.

Following the training process, the YOLOv8s model was evaluated on the unseen testing set to assess its effectiveness in tooth detection. The model's output consisted of detection results for each tooth, represented by bounding boxes that precisely define the location of each tooth within the image. Additionally, the model assigned corresponding tooth labels based on the Universal Numbering System, enabling accurate identification of tooth types. These detection results were then saved to a text file for further analysis.

KENNEDY CLASSIFICATION AND MISSING TEETH HANDLING

In this final stage, the results from the tooth detection phase were utilized to classify partially edentulous areas according to the Kennedy classification system, while also incorporating a mechanism to handle missing teeth. The whole process for tooth detection and Kennedy classification is shown in Figure 2.

The next step in our pipeline focused on identifying missing teeth which is crucial for subsequent Kennedy classification and potential modifications. This process involved comparing the teeth detected by the YOLOv8s model, stored in the text file (with information like bounding box coordinates and class labels), with the complete set of teeth numbers (#1 to #32). Any teeth absent from the YOLO model's detection were classified as missing.

ejprd.org - Published by Dennis Barber Journals.

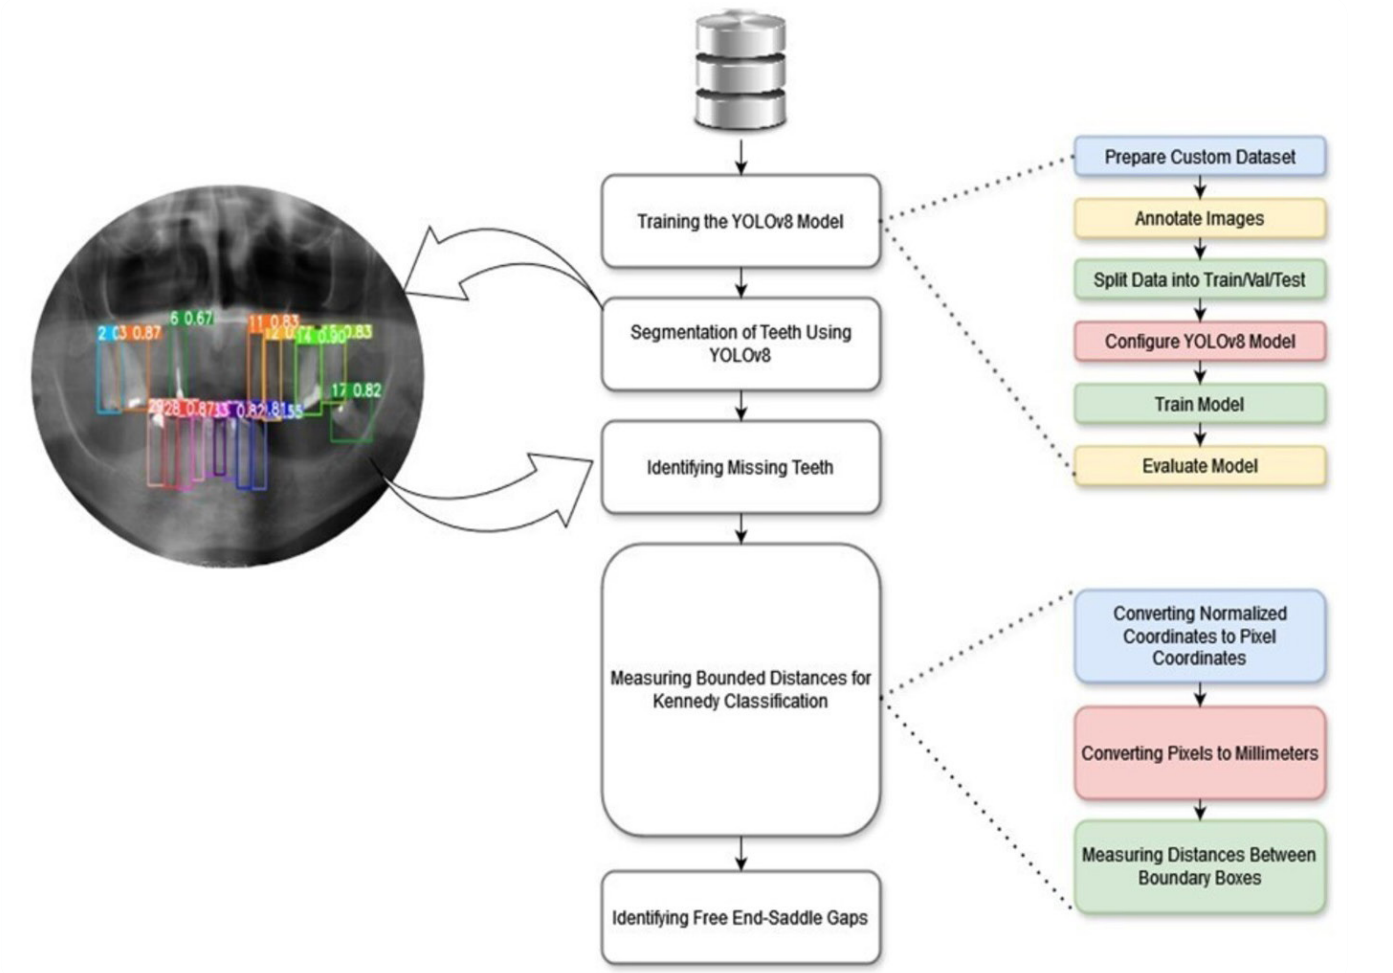

Figure 2: The proposed model for Kennedy Classification using YOLOv8s.

The measurement of bounded distances for Kennedy classification involves several key steps. First, for each detected tooth, the normalized coordinates obtained from the YOLO model's output—typically normalized between 0 and 1—are converted to pixel coordinates using the formula:

Pixel Coordinate = Normalized Coordinate × Image Dimension

Xpixel = x × W

Ypixel = y × H

where "x" and "y" are the normalized coordinates, and "W" and "H" are the width and height of the image in pixels, respectively.

Following this, a custom function named pixels_to_mm (pixels, dpi) is used to convert these pixel coordinates into millimetres. This function first converts pixels to inches by dividing by the image resolution (DPI) and then converts inches to millimetres using a multiplication factor of 25.4.

Once the teeth are localized in millimetres, the next step is to measure the distances between them to aid in Kennedy classification. This includes calculating the vertical line distance between the vertical edges of bounding boxes for neighbouring teeth bordering the missing teeth region. Distances less than a predefined threshold of 2 mm are considered insignificant and are therefore disregarded. Additionally, the centre-to-centre distance between the bounding boxes of neighbouring teeth is measured using a distance formula. Distances smaller than the sum of half the widths of the two bounding boxes are ignored, indicating no significant gap between the teeth.

In the final step, the analysis shifted to identifying free-end saddle gaps. This involved examining the list of detected teeth to locate consecutive missing teeth adjacent to only one neighbouring tooth. When such gaps were identified, they were classified as free-end saddle gaps, indicative of Kennedy Class I or Class II classifications. If no such gaps were detected, the dental arch was categorized as not fitting into any specific Kennedy class. Ultimately, the final Kennedy classification was determined based on the presence or absence of these identified free-end saddle gaps within the broader Kennedy classification framework. *(Figure 3)*

DATA ANALYSIS

The Precision, Recall, F1 score and mean average precision (mAP) were assessed to evaluate the detection accuracy of the trained model.

Precision: This metric focuses on the proportion of true positive predictions among all the positive predictions made by the model. A high precision value (closer to 1) indicates the model is less prone to false positives. In simpler terms, the model is more precise in identifying relevant instances and avoids incorrectly classifying irrelevant ones as positive.[8]

```
Input: Detected teeth list
Output: Kennedy classification
    Scan detected teeth list to identify consecutive missing teeth adjacent to only one neighboring tooth
if free-end saddle gaps identified then
        Classify as Kennedy Class I or Class II
end if
if no free-end saddle gaps identified then
        Classify as not corresponding to any Kennedy class
end if
```

Figure 3: Kennedy classification Algorithm.

$$Precision = \frac{Tp}{(Tp + Fp)}$$

Recall: This metric emphasizes the model's ability to correctly identify all actual positive cases. It essentially measures the percentage of true positive predictions out of all the actual positive instances present in the data. A high recall value signifies the model's effectiveness in capturing all the relevant cases and minimizes false negatives (missing actual positive instances).[8]

$$Recall = \frac{Tp}{(Tp + Tn)}$$

F1 Score: It provides a balanced view by incorporating both precision and recall into a single metric. The F1 score is calculated as the harmonic mean of precision and recall. A score closer to 1 signifies a well-performing model that strikes a good balance between identifying relevant instances accurately (precision) and capturing most of the actual positive cases (recall).[8]

$$F1\ Score = 2 * \frac{Precision * Recall}{Precision + Recall}$$

Mean Average Precision (mAP): For object detection tasks specifically, Mean Average Precision (mAP) is a crucial metric. It goes beyond simple classification (positive or negative) and considers the model's capability to accurately localize objects within an image. mAP takes into account both the precision and recall of the model across all object classes present in the dataset. A higher mAP value indicates the model's proficiency in not only identifying objects correctly (precision) but also pinpointing their exact locations within the images.[8]

VISUALIZATION AND EXPORTATION OF MODEL PREDICTION RESULTS

The visualization of model prediction results entailed presenting panoramic radiographic images annotated with distinctively coloured bounding boxes around each tooth. Each tooth was labelled with its corresponding number, and the Kennedy classification for each dental arch was displayed below the image. *(Figure 8)*

The process of exporting model prediction results involves generating a PDF report from annotated panoramic radiographic images. Each report includes detailed annotations such as missing and existing teeth, Kennedy classification, and a rough measurement of the edentulous area in millimetres. Therefore, each panoramic image was accompanied by a visual representation with annotated bounding boxes and numerical identifiers, textual descriptions under specific headings (e.g., Missing Teeth, Existing Teeth, Kennedy Classification) and a summarized measurement of the edentulous area to provide comprehensive clinical information.

RESULTS

TEETH DETECTION

The YOLOv8s model achieved high level of performance on the testing dataset, reaching a mean Average Precision (mAP50) of 98.1%. The precision and recall values were 95.7% and 95.8% respectively. *(Figures 4-6)*

The model achieved a training loss of 2.18 and a validation loss of 2.82. These loss values indicate the model's ability to learn from the training data while generalizing well to unseen data during validation.

Kennedy Classification with Class and Modification Areas

The model's effectiveness was evaluated in its ability to classify partially edentulous arches according to the Kennedy classification system, incorporating both class and modification areas. The results are presented in Table 1. Overall, the model achieved high performance across all metrics for both maxillary and mandibular arches. For the maxillary arches, the model achieved a precision of 0.946, recall of 0.897, and an F1-score of 0.909. For the mandibular arches, the precision was 0.977, recall was 0.966, and the F1-score was 0.968. The combined performance across both arches yielded a precision of 0.962, recall of 0.931, and an F1-score of 0.939. *(Figure 7)*

DISCUSSION

The current study had two primary objectives. Firstly, it aimed to achieve automated classification of partially edentulous arches according to the Kennedy system supplemented by Applegate's rules. Secondly, the model was designed to generate comprehensive reports that include the Kennedy classification (with any present modifications), the number of missing teeth, and type of the edentulous areas (whether free-end saddle, or bounded ones).

The model achieved a mean Average Precision (mAP50) of 98.1%, indicating high accuracy in tooth detection. The high precision (95.7%) and recall (95.8%) further underscore the model's reliability in identifying teeth while minimizing false positives and negatives, which is crucial for clinical applications.

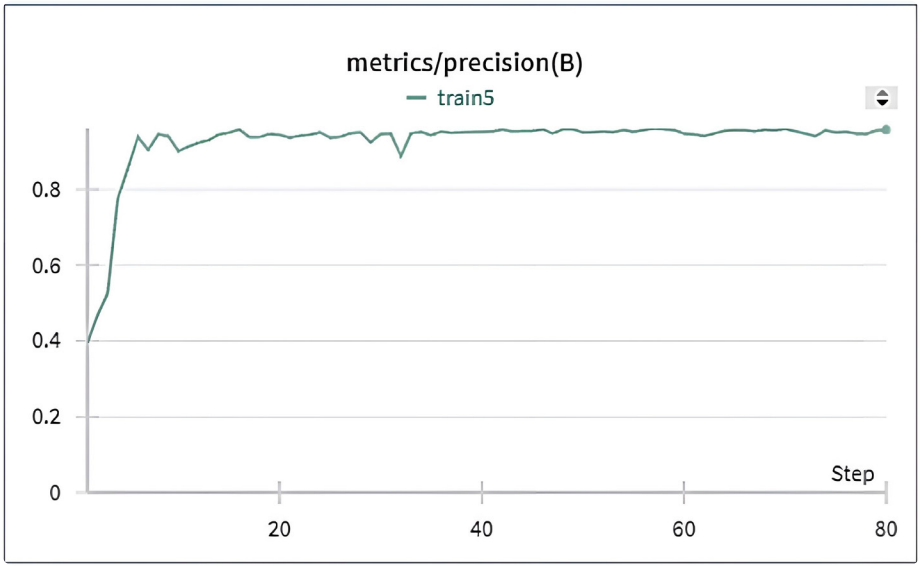

Figure 4: Precision of the teeth detection model over training steps. The graph shows the precision metric plotted against training steps, stabilizing above 0.8 after an initial rise, with minor fluctuations throughout the training process.

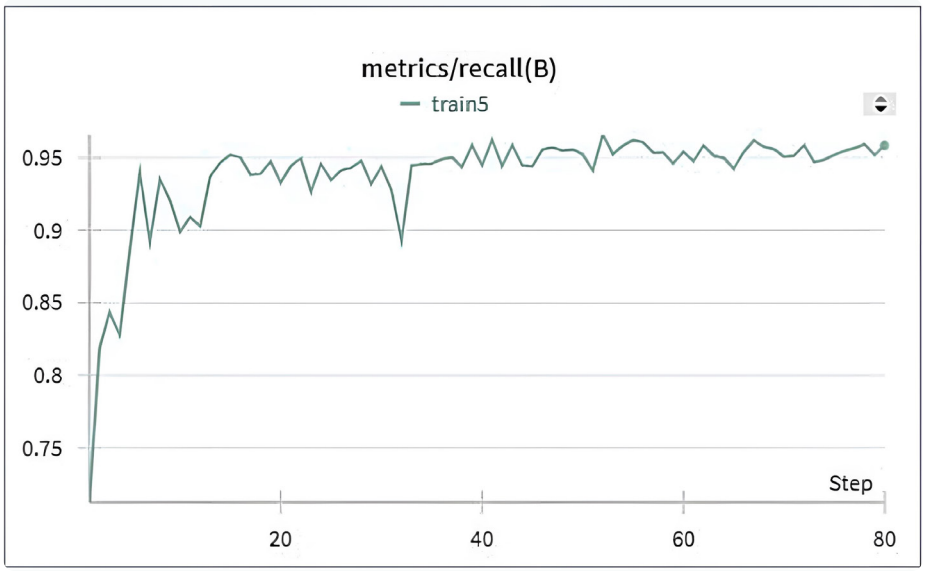

Figure 5: Recall of teeth detection model over training steps. The graph shows the recall metric stabilizing around 0.95 after initial fluctuations, with occasional minor drops throughout the training process, indicating high and consistent recall performance

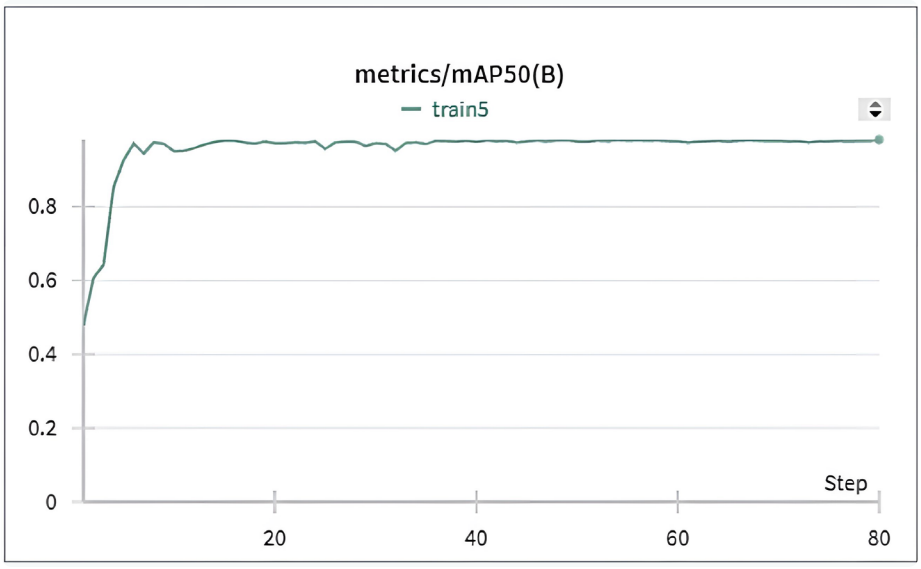

Figure 6: Mean Average Precision at IoU 50 (mAP50) of teeth detection model over training steps. The plot demonstrates that the mAP50 metric stabilizes above 0.85 after an initial rise, showing consistent performance with minor variations across the training steps.

Table 1. Model Accuracy in Kennedy classification.

Metric	Maxillary arch	Mandibular arch	Combined arches
Precision	0.946	0.977	0.962
Recall	0.897	0.966	0.931
F1 Score	0.909	0.968	0.939

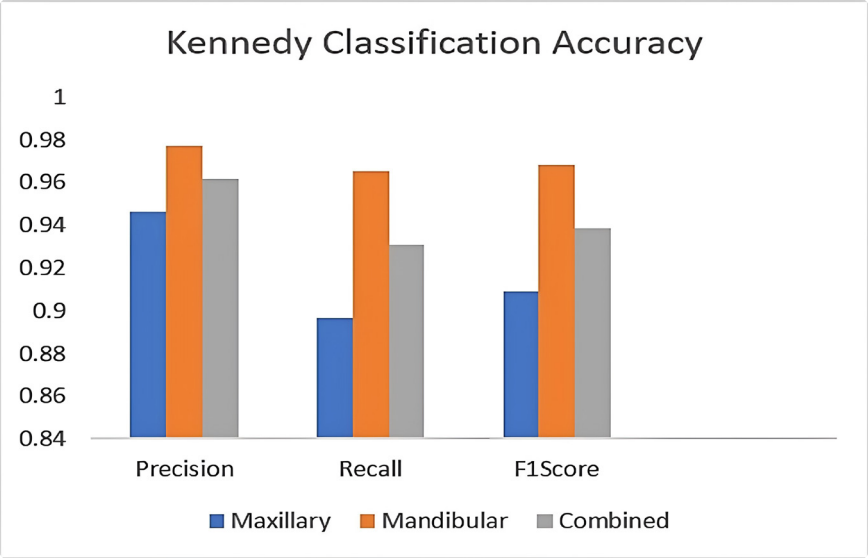

Figure 7: Model accuracy for Kennedy classification from panoramic radiographs.

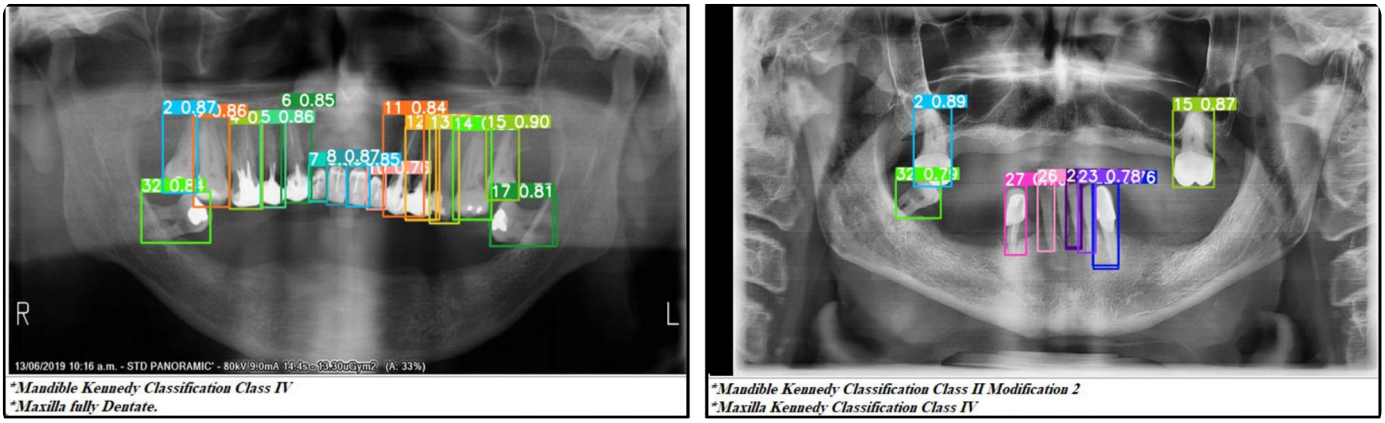

Figure 8: Screenshots of output results reporting the name of the identified Kennedy's class for the maxillary and mandibular arches, along with teeth identification.

The Kennedy classification results also highlight the robustness of the model. For maxillary arches, the model achieved an F1-score of 0.909, while for mandibular arches, the F1-score was 0.968. These metrics reflect the model's ability to balance precision and recall in accurately identifying and classifying partially edentulous areas. The slightly higher performance observed in mandibular arches could be attributed to differences in anatomical complexity or dataset representation. Future studies could explore these differences in greater detail to optimize performance across both arches.

The training loss (2.18) and validation loss (2.82) indicate the model's ability to generalize well without overfitting. This is supported by the effective use of data augmentation and validation during training, which ensured robust learning across diverse variations in panoramic radiographs. Compared to other studies employing different architectures, such as Faster R-CNN or RetinaNet, YOLOv8 demonstrates a superior balance between computational efficiency and detection accuracy, particularly in dental imaging tasks.[18]

In calculating vertical line distances between bounding boxes for teeth bordering missing regions, a predefined threshold of 2 mm was chosen based on considerations of biological relevance. Natural variations in tooth size and minor misalignments can lead to small gaps between teeth. Setting the threshold at 2 mm allowed us to focus on distances that are more likely to represent clinically significant spaces important for Kennedy classification and the relevant modification areas, such as those arising from missing teeth. While the 2 mm threshold appears effective in our study, further investigations could explore the impact of different thresholds on classification accuracy, particularly for datasets with highly variable tooth sizes or complex occlusal patterns.

Our decision to utilize the YOLOv8 deep learning architecture for this study was driven by several key factors. YOLO stands out as a formidable competitor to both two-stage detectors and earlier single-stage detectors, excelling in both accuracy and inference speed. Its popularity in practical applications stems from its straightforward architectural design, minimal complexity, and ease of implementation, making it a preferred choice for real-world deployment.[19] Unlike traditional two-stage detectors such as Faster R-CNN, which first generate region proposals and then classify them, YOLO performs object detection in a single pass. This efficiency is particularly advantageous for real-time applications and large datasets like ours, where computational demands can become a bottleneck. Additionally, YOLOv8 incorporates advancements such as a more powerful backbone (CSPDarknet53) and a feature pyramid network (neck) that enriches feature representation across scales.[17] These improvements enhance its ability to detect objects with varying sizes and complexities, which aligns well with the requirements of dental X-ray analysis.

Other neural networks, such as RetinaNet or Mask R-CNN, have also been employed for object detection in medical imaging, including dentistry. For example, Mask R-CNN is often used for segmentation tasks, providing pixel-level precision, which could be beneficial for certain applications like lesion detection.[20] However, segmentation models tend to be more computationally expensive and less suited for tasks focused on bounding box localization, such as Kennedy classification. Similarly, Faster R-CNN provides high detection accuracy but is slower than YOLO, making it less practical for clinical workflows that require efficient processing of large numbers of radiographs.[18]

In the context of dental applications, prior studies have highlighted YOLO's advantages. George *et al.* demonstrated its effectiveness in detecting dental features like cavities, implants, and impacted teeth from panoramic radiographs, achieving high accuracy with minimal computational overhead.[21] Similarly, Agulan *et al.* and Hassan *et al.* reported YOLO's superior performance in lesion detection and implant identification compared to other architectures.[8,22] These findings underscore YOLO's versatility and efficiency in handling dental imaging tasks, making it the preferred choice for this study.

Despite the promising results, certain limitations should be noted. First, the model was evaluated on publicly available datasets, which may not fully capture the diversity of clinical imaging scenarios. Variations in imaging equipment, patient demographics, and radiographic quality may influence real-world performance. Second, while the current study demonstrated high accuracy, additional testing on larger and more diverse datasets is necessary to confirm generalizability. Finally, the focus on bounding boxes rather than segmentation may limit certain clinical applications, such as precise tooth morphology analysis.

This study is among the first to use panoramic radiographs for automated classification of partially edentulous arches according to the Kennedy classification system with Applegate's rules. The study findings pave the way for streamlining diagnostic workflows in clinical dentistry. This approach could reduce manual effort, enhance classification accuracy, and enable standardized reporting across dental practices.

CONCLUSION

The YOLOv8 model provides an accurate and efficient solution for detecting teeth and classifying partially edentulous arches according to the Kennedy system and Applegate's rules using panoramic radiographs. This novel approach has the potential to improve diagnostic workflows and promote the integration of AI systems into clinical dentistry. Future research should focus on expanding the dataset, exploring segmentation-based approaches, and investigating real-world implementation to further enhance the model's clinical utility.

REFERENCES

1. Agbor, A.M., Kaptue B., Salomon, Z., Ananfack, G.C. and Tchweguem, C. Pattern and consequences of non-replacement of missing teeth amongst edentulous adults in ngaoundéré-cameroon. *British Journal of Healthcare and Medical Research. 2022;* **9**(3).
2. Skinner, C.N. A classification of removable partial dentures based upon the principles of anatomy and physiology. *J Prosthet Dent.* 1959; **9**:240-246.
3. Miller, E.L. Systems for classifying partially dentulous arches. *J Prosthet Dent.* 1970; **24**:25-40.
4. Henderson, D.W. and Le P. Langlois, S. *The Lung: Radiologic-Paihologic Correlations| E. Robert Heitzman, The Lung: Radiologic-Paihologic Correlations,* 2nd, CV Mosby Company, St. Louis, Toronto, Princeton (1984), p. 546,(Aust. Distr. CIG Medishield, North Ryde, NSW) ISBN 0-8016-2135-6. illustrated. (1985): 684.
5. Şakar, O. Classification of partially edentulous arches. *In Removable Partial Dentures: A Practitioners' Manual,* pp. 21-25. Cham: Springer International Publishing, 2024.
6. Şakar, O. Diagnosis and Treatment Planning in Partially Edentulous Patients. *In Removable Partial Dentures: A Practitioners' Manual,* pp. 41-56. Cham: Springer International Publishing, 2024.

7. Khanagar, S.B., Alfouzan, K., Awawdeh, M., Alkadi, L., Albalawi, F. and Alfadley, A. Application and performance of artificial intelligence technology in detection, diagnosis and prediction of dental caries (DC)—a systematic review. *Diagnostics (Basel).* 2022; **12**:1083.
8. Hassan, N.A., Kamel, A.E., Omran, A.E., Gad, M.W., Ashraf, N.M., Ahmed, O.M. and Abdel-Fattah, M.A. Automated identification of dental implants: A new, fast and accurate artificial intelligence system. *Eur J Prosthodontic Restor Dent.* 2024; **32**:162-167.
9. Khanagar, S.B., Alkadi, L., Alghilan, M.A. Kalagi, S. Awawdeh, M., Bijai, L.K. *et al.* Application and performance of artificial intelligence (AI) in oral cancer diagnosis and prediction using histopathological images: a systematic review. *Biomedicines.* 2023; **11**:1612.
10. Patil, S., Joda, T., Soffe, B., Awan, K.H., Fageeh, H.N., Tovani-Palone, M.R. *et al.* Efficacy of artificial intelligence in the detection of periodontal bone loss and classification of periodontal diseases: A systematic review. *J Am Dent Assoc.* 2023; **154**:795-804.e1.
11. Hakan, A., Jaju, P.P., Ezhov, M., Gusarev, M., Atakan, C., Sanders, A., *et al.* Development and validation of an artificial intelligence software for periodontal bone loss in panoramic imaging. *Int J Imaging Syst Technol.* 2024; **34**:e22973.
12. Mihai, T., Zhou, Y., Antonelli, A. and Becker, K. The application of artificial intelligence for tooth segmentation in CBCT images: A systematic review. *Appl Sci* 2024; **14**:6298.
13. It, S.C.D. Dental X-rays. (2023). https://www. kaggle.com/datasets/shanecandoit/dental-xrays.
14. Satishwar, I. (2023). Dental Panoramic Scan. figshare. Dataset. https://doi.org/10.6084/m9.figshare.22580263.v2.
15. Linear, T.I.N. Dentex Challenge 2023. https://www.kaggle.com/datasets/truthisneverlinear/dentex-challenge-2023.
16. Peck, S. and Peck, L. Tooth numbering progress. *The Angle Orthodontist.* 1996: **66**:83-84.
17. Terven, J., Córdova-Esparza, D-M. and Romero-González, J-A. A comprehensive review of yolo architectures in computer vision: From yolov1 to yolov8 and yolo-nas. *Machine Learning and Knowledge Extraction 5.* 2023; **4**:1680-1716.
18. Tan, L., Huangfu, T., Wu, L., *et al.* Comparison of YOLO v3, Faster RCNN and SSD for Real-Time Pill Identification. 30 July 2021, PREPRINT (Version 1) available at Research Square [https://doi.org/10.21203/rs.3.rs-668895/v1]
19. Diwan, T., Anirudh, G. and Tembhurne, J.V. Object detection using YOLO: Challenges, architectural successors, datasets and applications. *Multimed Tools Appl.* 2023; **82**:9243-9275.
20. Huang, S-T., Chu, Y-U., Liu, L.R., Yao, W.T., Chen, Y.F., Yu, C.M. *et al.* Deep learning-based clinical wound image analysis using a mask R-CNN architecture. *Journal of Medical and Biological Engineering.* 2023; **43**:417-426.
21. George, J., Hemanth, T. S. Raju, J., Mattapallil, J.G. and Naveen, N. *Dental radiography analysis and diagnosis using YOLOv8.* In 2023 9th International Conference on Smart Computing and Communications (ICSCC), pp. 102-107. IEEE, 2023.
22. Agulan, R. Z., Magsumbol, J-A.V. and Rosales, M.A. *Detection of dental caries based on GV black classification utilizing Yolov8 from X-ray images.* In 2024 IEEE International Conference on Imaging Systems and Techniques (IST), pp. 1-6. IEEE, 2024.