Tooth Detection and Numbering in Panoramic Radiographs Using YOLOv8-Based Approach

Felipe Rogério Silva Teles[1(✉)], Alison Corrêa Mendes[1], Anselmo Cardoso de Paiva[1], João Dallyson Sousa de Almeida[1], Geraldo Braz Junior[1], Aristófanes Corrêa Silva[1], and Pedro De Alcantara Dos Santos Neto[2]

[1] Núcleo de Computacão Aplicada - Universidade Federal do Maranhão (UFMA), Caixa Postal 65.085-580, São Luís, MA, Brazil
frs.teles@discente.ufma.br, {alison.mendes,paiva}@nca.ufma.br, {joao.dallyson,geraldo.braz,ac.silva}@ufma.br

[2] Departamento de Computação, Universidade Federal do Piauí, Teresina, Brazil
pasn@ufpi.edu.br

Abstract. Before a dental professional performs any procedure or diagnosis, they need to know the patient's dental arch. For that, it is common for them to ask the patient to take a panoramic radiograph. The use of neural networks to assist this professional in this stage is not recent, and most studies use segmentation networks to solve the problem. However, the segmentation result does not make explicit the specific position of the tooth and its numbering according to the international system (FDI), presenting only more specific details. In this study, we aimed to use a powerful and efficient detection neural network called You Only Look Once v8 to perform automated tooth detection and numbering based on FDI, using a dataset that contains 166 anonymized and deidentified panoramic dental radiographs of patients from Noor Medical Imaging Center, Qom, Iran, and are public. Labels were created using an online tool for production in the YOLO standard. The metrics used to evaluate the trained model were precision, recall, and mAP50. The results of each were 0.95818, 0.95505, and 0.97384. The conclusion of the study uses the model training generated a weight to test the model in a real-world scenario.

Keywords: tooth detection · panoramic radiography · deep learning · dental image

1 Introduction

One of the most important exams for dental professionals is the panoramic radiograph, and it's the most popular screening test prescribed in dentistry too [6]. It is also a util and necessary tool in several areas of dentistry, mainly in the

A. Cunha et al. (Eds.): MobiHealth 2023, LNICST 578, pp. 239–253, 2024.
https://doi.org/10.1007/978-3-031-60665-6_18

diagnosis and treatment of oral diseases such as caries, periodontal diseases, and oral pathologies [12,23]. It is noteworthy that in addition to being a low-cost exam, those who perform are not so exposed to radiation. Radiograph also provides images of the dental arch and the mandibular and maxillary region. They recommended that not all patients must be tested and those who do be limited to areas necessary for proper diagnosis and treatment based on the good exercise of professional judgment [26].

When the dentist has images of the examination, he can examine it and then diagnose it. To increase the accuracy of diagnosis, tooth type, and teeth numbering should be identified and determined based on their anatomy and location [5]. When discussing a small number of exams, it is possible to do the procedure manually and carefully. But in practice, a dentist has several patients, and the amount of exams he has to analyze is large [18], which causes more errors to occur because of fatigue and the fact that factors such as subjectivity and experience influence the analysis [16].

The use of computers and systems is increasing more and more in the health area, and dentistry is no different because it facilitates diagnosis and, consequently, simplifies treatment planning [4]. That being said, one of the main technologies that have been used is Artificial Intelligence (AI), which over time has become frequent in many areas. They make the computer mathematically imitate how the human brain thinks and makes decisions [25]. Recently, AIs have experienced accelerated development and have become one of the most influential innovations worldwide, and AI-based methods are beginning to be used to assist dentists in interpreting radiographic results [29]. These methods help to make the identification and classification faster and reduce the errors from fatigue. Moreover, the increase of AIs could provide fully automated dental chart filling and treatment planning, giving more time to the dentist to perform treatments [5].

Artificial intelligence is a broad area, so the focus of this study was on Deep Learning(DL), which is a computer-assisted framework and a subarea of Machine Learning (ML) [30]. Deep Learning (DL) is considered the most cost-effective and time-efficient ML approach [8]. An advantage of DL is processing large amounts of data, such as text, sound, and images [20]. One of the basic types of architecture that works well in DL is the Convolutional Neural Network (CNN). For color and boundary detection functions, CNNs prove very accurate in image detection. Besides, they have several emerging applications in medical image analysis due to the advent of computing hardware algorithms and expansion in the amount of data [4].

This study aims to carry out the automated detection of tooth numbering in panoramic radiography exams using the most modern version of the YOLO v8 model, which allows training with more quality and less data [9], to evaluate its performance using precision, recall, and mean average precision (map) as metrics. After obtaining the results, we hope it can serve as an alternative to assist dental professionals in detecting and numbering teeth accurately and quickly.

2 Related Works

Miscellaneous previous studies have used Deep Learning techniques for the task of numbering teeth in radiographs using CNN-based architecture and various methods. Al-Sarem et al. [2] used different networks to try to find the position and number of the missing teeth. For this task, they used a dataset with 120 images, and the best results for segmentation and classification were using, respectively, a U-Net and DenseNet169 pre-trained model. The best result for segmentation was an accuracy of 90.81%, and for classification, the precision was 0.88 and recall 0.98. Görürgöz et al. [10] was designed based on R-CNN inception architecture that automatically detect and number the teeth on periapical images. They used a dataset with 660 images, and the results of F1 score, precision, and sensitivity were 0.8720, 0.7812, and 0.9867, respectively. The study of Alam et al. [3] consists of segmenting teeth into two modules, the first module designed to identify teeth and the second module designed to number teeth. For tooth numbering, they trained the VGG16 convolution architecture using 1300 images and evaluated the model using 200 images (1500 in total). Results obtained for teeth diagnosis were 89.8% and, for teeth numbering, precision of 86.5%.

In dental radiology, You Only Look Once(YOLO) v4 has already been used in detecting tooth numbering. Putra et al. [20] used YoloV4 and trained a dataset that contained 500 images, 400 for training and 100 for testing, and reached a result of 88.5% accuracy and 87.70% precision. Astuti et al. [5] also used YOLOv4 with the same amount of images to train and validate the model. However, they used other metrics to evaluate the model, and they achieved a sensitivity of 99.42% and a specificity of 87.06%.

This study experiments with an update of these approaches using a more up-to-date convolutional neural network called YOLOv8, which is a continuation of You Only Look Once v4 and has more precision and less need for data, which directly affects processing [11]. The obtained results were promising and demonstrated that YOLOv8 applied to the problem of automatic tooth detection performs well.

3 You only Look once

You Only Look Once (YOLO) is a potent object detection algorithm that detects objects in real-time using CNN16. The YOLO system treats object detection as a single regression problem, which makes it possible to discover which classes are present and where they are [21]. Currently, YOLO has eight versions, and its effectiveness is evidenced through a comparative analysis carried out in Terven et al. [27] that proves its efficiency and makes a comparative analysis of each version released in their work.

The annotation format used by the YOLO model consists of five params, and there are, respectively, the class of the object, the X and Y of the bounding box center, and the width and height of the bounding box as seen below:

$$(ObjectId - class)(Xcenter)(Ycenter)(Width)(Height) \quad (1)$$

The initial convolutional layers of the network resize the image according to the chosen input size and extract features, while the fully connected layers predict the output probabilities and coordinates [21]. The input image is split into s × s grid cells (default = 7 × 7), with each cell predicting (B) bounding boxes, each containing five parameters and sharing prediction probabilities of classes (C) [11]. How appears in Fig. 1.

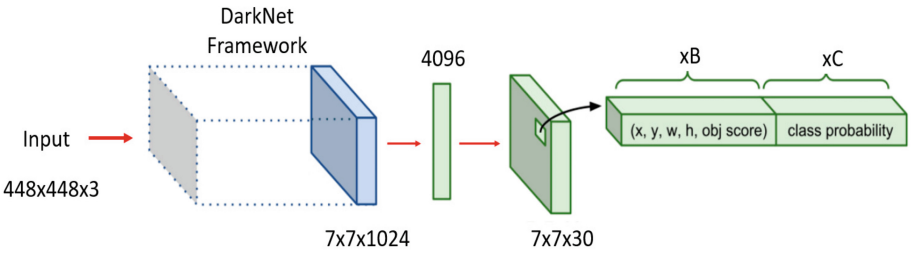

Fig. 1. General architecture of YOLO using a input size of 448 × 448 and 3 classes [11].

3.1 Bag of Freebies

Bag of Freebies (BoF) is the set of low computational cost techniques that YOLO has improved to improve YOLO training performance as they avoid overfitting and improve model generalization [32]. Examples of these techniques is Data augmentation.

Data Augmentation. Expanding the training dataset with variations of the original data using random geometry transformation and random color flickering [32]. This technique allows the YOLO model to learn efficiently with a much smaller number of images than a pure neural network would need [9].

3.2 YOLO V8

YOLO's latest version, YOLO v8, was officially released in January 2023 by Ultralytics, and even though it is considered incomplete about the level of functionality of previous versions, comparisons demonstrate its superiority as the new YOLO state-of-the-art[18]. The Fig. 2 shows a comparison between previous versions. YOLOv8, unlike YOLOv4, features an anchor-free system that allows data to be processed independently of objectivity, classification, and regression tasks. This design allows each branch to focus on its labor and improves the overall accuracy of the model [27]. Furthermore, YOLOv8 uses CIoU [33] and DFL [15] loss functions for bounding box loss and binary cross-entropy for classification that directly leads to improved object detection performance, especially when dealing with smaller objects [27].

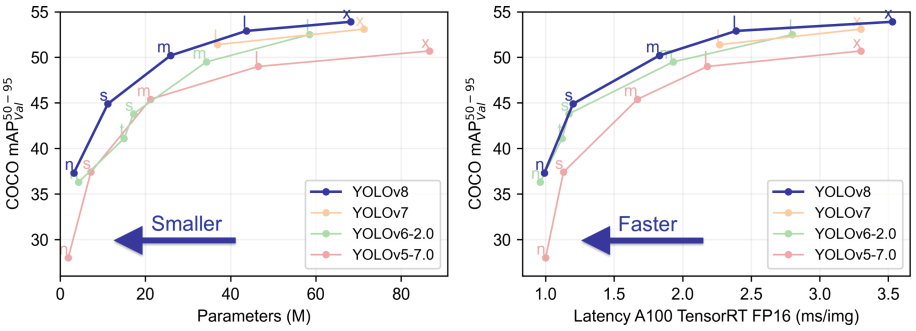

Fig. 2. Comparison chart between latest versions of yolo. [13]

4 Matherial and Methods

4.1 Tooth Numering

Tooth numbering is the most commonly used practice in dentistry to identify specific teeth in a patient's mouth and is an excellent way to perform a pre-diagnosis.

The Fédération Dentaire Internationale (FDI) Tooth Numbering System, officially developed by the World Dental Federation [19], divides the mouth into four quadrants: Upper Right (UR), Upper Left (UL), Lower Left (LL), Lower Right (LR), and divides each quadrant into eight teeth: Central Incisor (1), Lateral Incisor (2), Canine (3), First Premolar (4), Second Premolar (5), First Molar (6), Second Molar (7), Third Molar (8). In this way, the identification of the tooth occurs through two numbers, the first referring to the quadrant and the second referring to the tooth itself, as shown in Fig. 3.

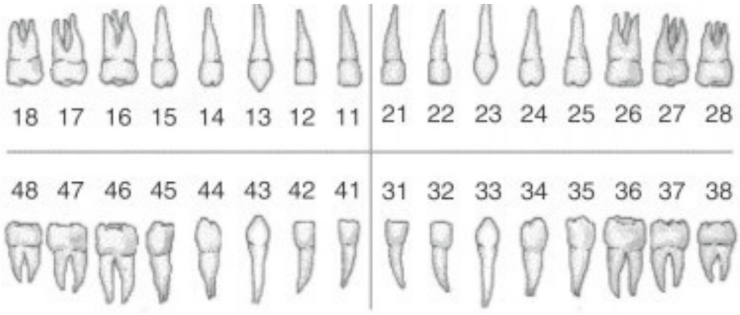

Fig. 3. Tooth numbering by FDI system [17]

This standardized numbering system simplifies communication between dental professionals and ensures that everyone, regardless of teeth location, can easily understand and reference specific teeth [28]. The practitioner needs to be

familiar with the FDI Tooth Numbering System to document and communicate dental conditions and treatment plans accurately.

4.2 Dataset

The data obtained is part of a dataset containing almost 2,000 panoramic radiographs taken by the Soredex CranexD [1] digital panoramic X-ray unit. For more accurate detection, selected exams of individuals over 20 years old because the jaw growth was almost complete and they no longer had baby teeth [14]. In addition, we did not use images of patients with any implant kind.

This study used a dataset with, in total, 166 images from the original dataset, where 150 and their respective labels were used in the stage of training and validation of the model, and 16 non-labeled images were used to test the trained model in a real-world context.

4.3 Labeling Dataset

The labels were obtained manually through the free-to-use tool MakeSense (SkalskiP from GitHub) [24], and each tooth represented its respective tooth numbering class. The classes used following the numbering indicated in the international numbering system was 32 [19] and how it showed in the Fig. 3. There is an example of marking the bounding box in Fig. 4.

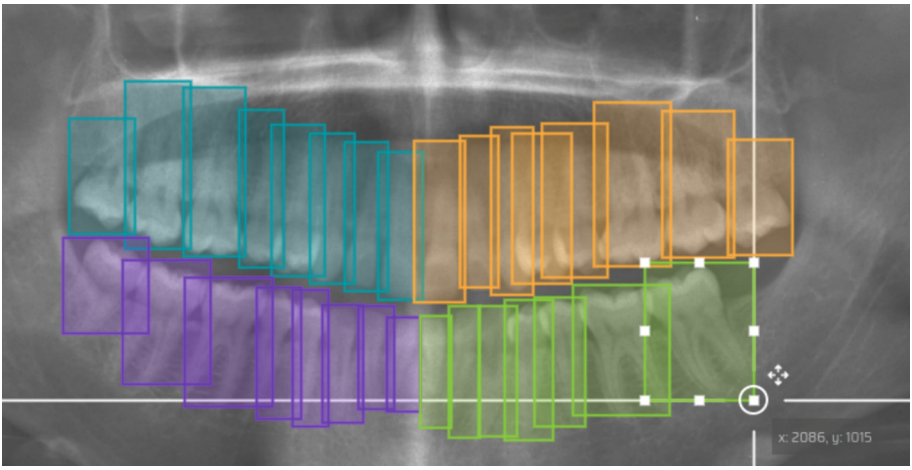

Fig. 4. Example image labeling on MakeSense tool that the four quadrants labels separated by color.

The total number of edentulous areas found was 4210, and the quantity of each one is present at Table 1. Not all patients had all their teeth. Therefore, the number of edentulous areas was not equal for all classes.

Table 1. On the right of the table you can see the amount of the teeth in the mandible region, and on the left you can see the amount of the teeth in the inferior jaw.

Maxillary Tooth	Labels Amount	Mandibular Tooth	Labels Amount
11	148	31	150
12	144	32	147
13	150	33	147
14	136	34	145
15	128	35	131
16	132	36	121
17	140	37	134
18	79	38	86
21	149	41	147
22	145	42	145
23	141	43	142
24	135	44	144
25	126	45	134
26	136	46	116
27	134	47	128
28	80	48	90

4.4 Training

With the notes completed, the training process begins. To carry out the study, the cloud development environment, Google Collaboratory, also called Google Colab, was used, which provides functional resources such as the Graphics Processing Unit (GPU), which reduces training time [7]. Google Colab proved efficient in training and saving memory on the local machine, as it allows integration with Google Drive. The model chosen for the train was YOLOv8x [13]. For training, 150 images of varying sizes were used. As seen in Sect. 3, the YOLO architecture resizes the image according to the size chosen in the training input. After a series of training attempts, it was concluded that the best input size for the problem was 640×640, and due to the low amount of data, training was carried out with batch size 6, requiring 100 epochs to reach the presented result.

4.5 Evaluating the Model

To evaluate the model presented, the metrics used were accuracy, precision, and mean Average Precision (mAP). However, before explaining each of the metrics, it is necessary to explain some basic concepts of variables used in the evaluation calculation of results, namely False Positives (FP), False Negatives (FN), and True Positives (TP).

False positives (FP) are the objects that the model detected wrongly. So, there are cases in which the model predicted an object, but there is no

corresponding truth object. False Negatives (FN) are objects that the detection model was unable to detect despite representing instances where there are real objects in the image, which means that the model did not predict the bounding boxes accurately. True Positives (TP) are the objects that the model correctly detected. In other words, TP represents cases where the model's predicted bounding boxes almost right overlap the ground truth bounding boxes and are classified correctly.

Recall. Is the fraction of relevant instances that were retrieved, and it can be calculated using the following operation:

$$Recall = \frac{TP}{TP + FN} \tag{2}$$

Precision. As a result of calculating the precision, a measure of how many of the positive predictions made are correct is obtained, and it can be represented as follows:

$$Precision = \frac{TP}{TP + FP} \tag{3}$$

Intersection over Union. In order to grasp the concept of Mean Average Precision, which is the principal metric used to evaluate the quality of a detection model, it is essential to have a clear understanding of Intersection over Union (IoU). IoU is the most popular evaluation metric used in object detection benchmarks [22]. The bounding box referring to the labels presented in the input is known as the ground truth. Calculating the IoU is necessary to determine the gap between the bounding box demarcated by the model and the ground truth.

Mean Average Precision. It is calculated by finding Average Precision(AP) for each class and then averaging over several classes. Average precision (AP) is how the precision recovery curve is summarized as a single value representing the average of all precisions. Below is the Mean Average Precision calculation:

$$mAP = \frac{1}{n} \sum_{i=1}^{n} AP_i \tag{4}$$

For this study, map50 and map95 were used, and they refer to two different mAP calculations based on different IoU thresholds. The mAP50 uses an IoU threshold of 0.5 and is the most important evaluation indicator [31]. While mAP95 is using an IoU threshold of 0.95.

4.6 Testing

To test the model, we chose to use the 16 images that do not have a label, and we used the prediction function based on the weights obtained after training. It

consists of the YOLO model analyzing the input image using its learned parameters and producing bounding boxes around the detected objects, corresponding class labels, and confidence scores. Contributing directly to evaluating the model in a practical scenario, that is, in the real world.

5 Results

Several training sessions were carried out to achieve the most promising result, carried out with 120 images and validated with 30. After one hundred epochs in a batch of six, we arrived at the results that will be presented in this session.

5.1 Confusion Matrix

It's workable to see the model's performance in a confusion matrix. The confusion matrix is a table with the number of rows and columns varying according to the number of classes and which reports the number of false positives, false negatives, true positives, and true negatives. It allows you to have more detailed information than just calculating metrics. As our experiment has many classes, the confusion matrix focuses on the diagonal, the number of negative verities we obtained. The plotted confusion matrix is in the Fig. 5.

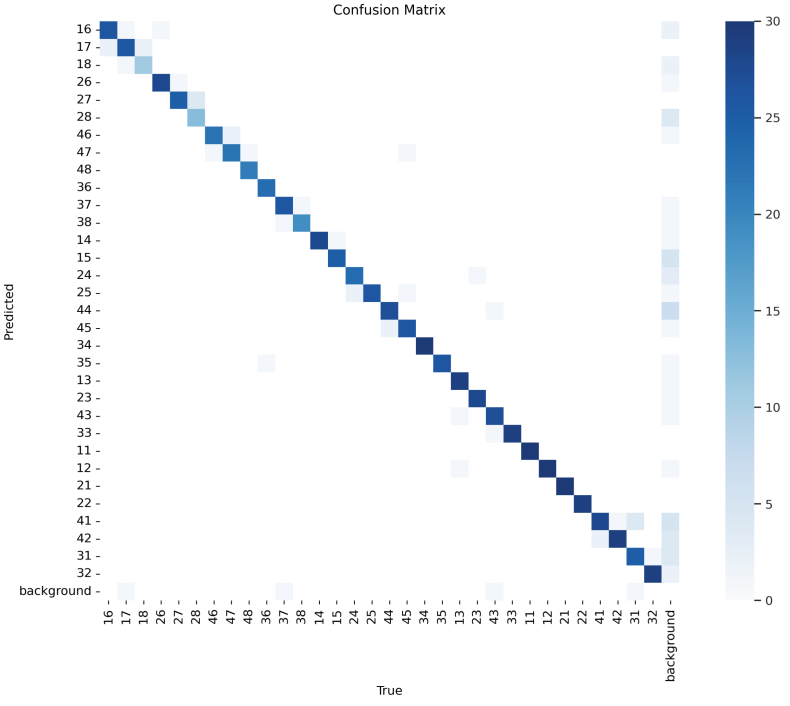

Fig. 5. Quantitative indicator of TP and FP for each class in a confusion matrix.

5.2 Qualitative Results

After knowing the number of false negatives, false positives, true negatives, and true positives, YOLO automatically generates the result in a CSV format, where each column has a metric, and each row represents an epoch. The best results obtained at different times and the precision, recall, and mAP50 were respectively 0.95818, 0.95505, and 0.97384. In addition to the CSV file with results, we also have graphs relating the times and the results obtained, as shown in Fig. 6. Map95 was also used, and the best result was 0.67214, but it is less relevant than map50 [31].

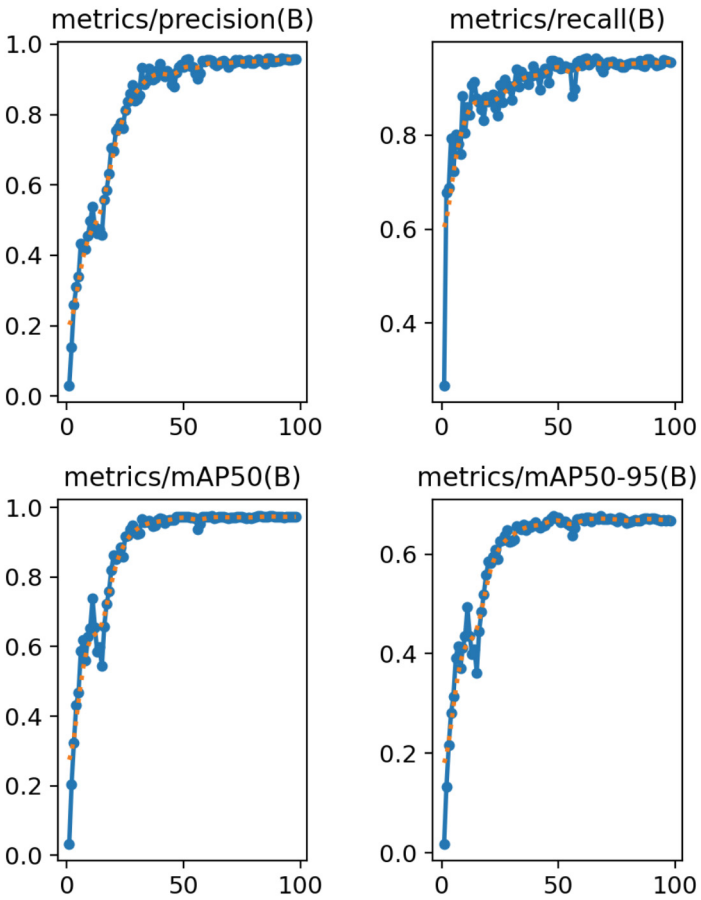

Fig. 6. The graph on the upper left side analyzes the precision of each epoch. The graph on the upper right side analyzes the recall per epoch. The graph on the lower left side analyzes the map with the IoU limited to 0.5 in each epoch. The last graph analyzes the map with the IoU limited to 0.95. Each metric present in the graphs is explained more clearly in Sect. 4.5

Analysis of the graphs shows stability in the main metrics from epoch 50 onwards, which proves the robustness of the model and the training process. This phenomenon indicates that the model has converged to a stable solution and that increasing the number of epochs is not a guarantee of improvement in the model's performance.

This stability is usual in deep learning models. It indicates that the model has learned patterns in the data well and has reached a point where additional training does not improve model performance. Therefore, it is worth a thorough analysis to ensure that the model hits a good level of generalization, that is, the model's ability to recognize tooth patterns learned from radiography images that were not present during training.

In summary, stabilization suggests that the model has reached a good level of training, and it is prepared for validation and testing to evaluate its concrete generalization performance in tooth numbering.

5.3 Quantitative Results

To carry out a more precise analysis, we evaluated the individual performance of each class according to tooth number. This step helps to identify whether the number of instances of each class influences the result. The application of metrics to each tooth class is shown in Table 2

Table 2. On the right of the table, you can see the results of the teeth in the mandible region, and on the left, you can see the results of the teeth in the inferior jaw.

Maxillary Tooth	Precision	Recall	Mandibular Tooth	Precision	Recall
11	0.981	1	31	0.962	0.846
12	0.975	1	32	0.967	0.983
13	0.994	0.968	33	0.953	1
14	0.965	0.976	34	0.989	1
15	0.918	1	35	0.913	1
16	0.943	0.959	36	1	0.985
17	0.881	0.964	37	0.927	0.91
18	0.794	0.862	38	0.907	0.979
21	0.988	1	41	0.829	0.97
22	0.979	1	42	0.941	0.967
23	0.985	1	43	0.993	0.933
24	0.94	1	44	0.977	0.966
25	0.899	1	45	0.937	0.929
26	0.966	0.989	46	0.881	0.913
27	0.859	0.939	47	0.867	0.917
28	0.884	0.895	48	0.951	0.909

With the individual analysis of each class, we came to the conclusion that the trained model presented a satisfactory result in the tooth numbering and detection task and is ready to perform the test. To check the model's generalization performance, we separated 16 images not present in the training stage and applied the generated weights.

5.4 Testing the Model

At the end of the training, the model-generated weights were used in this study to evaluate the model in a real-world application, that is, in a practical way and with images unknown to the model. This step contributes to automating the tooth detection and numbering process, as new training is unnecessary.

To test the model, we used 16 panoramic radiograph images that were not present in the training to evaluate the model in a real-use situation in an automated way and how the model behaves with generalization. In the Fig. 7, one of the images used to test the trained model.

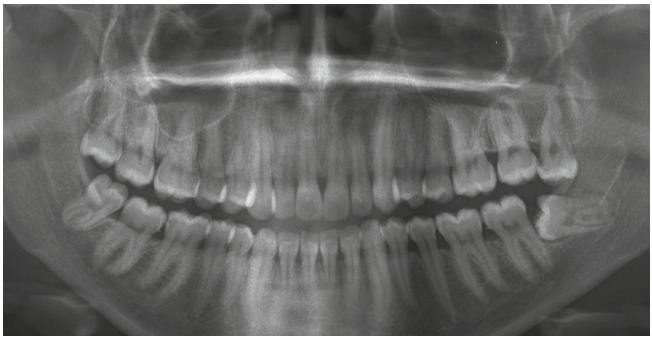

Fig. 7. Example of an image used to test the prediction of the model.

The prediction mode allows us to use the weights generated during training to carry out detection automatically, that is, without having to train the model again. Figure 8 has an example of successfully performed multi-object detection on panoramic radiographs with complete dentition and an edentulous area by YOLO v8.

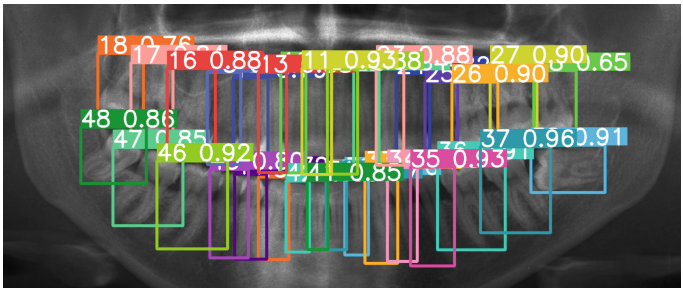

Fig. 8. Result of tooth numbering predict using You Only Look Once v8 on panoramic radiograph.

6 Conclusion

This work used a modern detection neural network based on deep learning methods to automatically execute the task of detecting each of the 32 teeth in panoramic radiographic images. To develop the model proposed in this study, we used 100 teeth images and extracted 2718 teeth, implant, and crown data. The model can number each tooth. The experimental results showed a high accuracy in tooth numbering using YOLO v8. As a result, the precision, recall, and mAP50 were respectively 0.95818, 0.95505, and 0.9738.

Through the weights obtained with the training, it is possible to perform the detection automatically using the prediction method of YOLO. It also allows the weights to carry out new models in future studies. Therefore, we consider the model presented as scalable.

Finally, it is expected that this system be used as decision support to help dentists as a way of helping them to save diagnosis time, using in practice only one computer in the clinic to improve the oral health system and increase the capacity of people that a professional can meet, and contribute to the development of radiology-related algorithms.

Acknowledgements. This work was carried out with the support of the Coordination for the Improvement of Higher Education Personnel - Brazil (CAPES) - Financing Code 001, Maranhão Research Support Foundation (FAPEMA), National Council for Scientific and Technological Development (CNPq) and Brazilian Company of Hospital Services (Ebserh) Brazil (Proc. 409593/2021-4).

The results presented in this paper were obtained through research carried out at the "CENTRO DE REFERÊNCIA EM INTELIGÊNCIA ARTIFICIAL - CEREIA", based at the Federal University of Ceará in partnership with the Hapvida NotreDame Intermédica Group and supported by the Grant 2020/09706-7, São Paulo Research Foundation (FAPESP).

References

1. Abdi, A.H., Kasaei, S., Mehdizadeh, M.: Automatic segmentation of mandible in panoramic x-ray. J. Med. Imaging **2**(4), 044003–044003 (2015)
2. Al-Sarem, M., Al-Asali, M., Alqutaibi, A.Y., Saeed, F.: Enhanced tooth region detection using pretrained deep learning models. Int. J. Environ. Res. Public Health **19**(22), 15414 (2022)
3. Alam, M.K., et al.: Teeth segmentation by optical radiographic images using VGG-16 deep learning convolution architecture with R-CNN network approach for biomedical sensing applications. Opt. Quant. Electron. **55**(9), 808 (2023)
4. Almalki, A., Latecki, L.J.: Self-supervised learning with masked image modeling for teeth numbering, detection of dental restorations, and instance segmentation in dental panoramic radiographs. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 5594–5603 (2023)
5. Astuti, E.R., et al.: The sensitivity and specificity of YOLO v4 for tooth detection on panoramic radiographs. J. Int. Dent. Med. Res. **16**(1), 442–446 (2023)
6. Bekiroglu, N., Mete, S., Ozbay, G., Yalcinkaya, S., Kargul, B.: Evaluation of panoramic radiographs taken from 1,056 Turkish children. Niger. J. Clin. Pract. **18**(1), 8–12 (2015)
7. Bisong, E., Bisong, E.: Google Colaboratory. Building Machine Learning and Deep Learning Models on Google Cloud Platform: A Comprehensive Guide for Beginners, pp. 59–64 (2019)
8. Dargan, S., Kumar, M., Ayyagari, M.R., Kumar, G.: A survey of deep learning and its applications: a new paradigm to machine learning. Arch. Comput. Methods Eng. **27**, 1071–1092 (2020)
9. Fabrice, N., Lee, J.J., et al.: SMD detection and classification using YOLO network based on robust data preprocessing and augmentation techniques. J. Multimedia Inf. Syst. **8**(4), 211–220 (2021)
10. Görürgöz, C., et al.: Performance of a convolutional neural network algorithm for tooth detection and numbering on periapical radiographs. Dentomaxillofacial Radiol. **51**(3), 20210246 (2022)
11. Hussain, M.: YOLO-v1 to YOLO-v8, the rise of YOLO and its complementary nature toward digital manufacturing and industrial defect detection. Machines **11**(7), 677 (2023)
12. Izzetti, R., Nisi, M., Aringhieri, G., Crocetti, L., Graziani, F., Nardi, C.: Basic knowledge and new advances in panoramic radiography imaging techniques: a narrative review on what dentists and radiologists should know. Appl. Sci. **11**(17), 7858 (2021)
13. Jocher, G., Chaurasia, A., Qiu, J.: YOLO by Ultralytics (2023). https://github.com/ultralytics/ultralytics
14. Koh, K., Tan, J., Nambiar, P., Ibrahim, N., Mutalik, S., Asif, M.K.: Age estimation from structural changes of teeth and buccal alveolar bone level. J. Forensic Leg. Med. **48**, 15–21 (2017)
15. Li, X., et al.: Generalized focal loss: learning qualified and distributed bounding boxes for dense object detection. In: Advances in Neural Information Processing Systems, vol. 33, pp. 21002–21012 (2020)
16. Lin, P., Huang, P., Huang, P., Hsu, H., Chen, C.: Teeth segmentation of dental periapical radiographs based on local singularity analysis. Comput. Methods Programs Biomed. **113**(2), 433–445 (2014)

17. Mourão, J., Sousa, J.: Lesão dentária na anestesiologia. Revista Brasileira de Anestesiologia **28** (2014). https://doi.org/10.1016/j.bjan.2013.04.009
18. Muresan, M.P., Barbura, A.R., Nedevschi, S.: Teeth detection and dental problem classification in panoramic X-ray images using deep learning and image processing techniques. In: 2020 IEEE 16th International Conference on Intelligent Computer Communication and Processing (ICCP), pp. 457–463. IEEE (2020)
19. Peck, S., Peck, L.: Tooth numbering progress. Angle Orthod. **66**(2), 83–84 (1996)
20. Putra, R.H., et al.: Automated permanent tooth detection and numbering on panoramic radiograph using deep learning approach. Oral Surg Oral Med Oral Pathol Oral Radiol **137**, 537–544 (2023)
21. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: unified, real-time object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 779–788 (2016)
22. Rezatofighi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I., Savarese, S.: Generalized intersection over union: a metric and a loss for bounding box regression. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 658–666 (2019)
23. on Scientific Affairs, A.D.A.C., et al.: The use of dental radiographs: update and recommendations. J. Am. Dent. Assoc. **137**(9), 1304–1312 (2006)
24. Skalski, P.: Make sense (2019). https://github.com/SkalskiP/make-sense/
25. Strong, A.: Applications of artificial intelligence & associated technologies. Science [ETEBMS-2016] **5**(6) (2016)
26. Sun, W., Xia, K., Tang, L., Liu, C., Zou, L., Liu, J.: Accuracy of panoramic radiography in diagnosing maxillary sinus-root relationship: a systematic review and meta-analysis. Angle Orthod. **88**(6), 819–829 (2018)
27. Terven, J., Cordova-Esparza, D.: A comprehensive review of YOLO: from YOLOv1 to YOLOv8 and beyond. arXiv preprint: arXiv:2304.00501 (2023)
28. Türp, J.C., Alt, K.W.: Designating teeth: the advantages of the FDI's two-digit system. Quintessence Int. **26**(7) (1995)
29. Tuzoff, D.V., et al.: Tooth detection and numbering in panoramic radiographs using convolutional neural networks. Dentomaxillofacial Radiol. **48**(4), 20180051 (2019)
30. Umer, F., Habib, S., Adnan, N.: Application of deep learning in teeth identification tasks on panoramic radiographs. Dentomaxillofacial Radiol. **51**(5), 20210504 (2022)
31. Zhang, S., Liu, J., Zhang, X.: Adaptive compressive sensing: an optimization method for pipeline magnetic flux leakage detection. Sustainability **15**(19), 14591 (2023)
32. Zhang, Z., He, T., Zhang, H., Zhang, Z., Xie, J., Li, M.: Bag of freebies for training object detection neural networks. arXiv preprint: arXiv:1902.04103 (2019)
33. Zheng, Z., Wang, P., Liu, W., Li, J., Ye, R., Ren, D.: Distance-IoU loss: faster and better learning for bounding box regression. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 34, pp. 12993–13000 (2020)