ISBN: 978-628-96613-3-0. ISSN: 2414-6390. Digital Object Identifier: https://dx.doi.org/10.18687/LACCEI2026.1.1.1953

Neural Network-Based Detection of Dental Caries Using Roboflow and Jetson Nano

Alicia María Reyes-Duke[1]; Silvio Javier Lázaro-Cárdenas;Aaron Ortíz-Pinedar[2]
[1,2]Universidad Tecnológica Centroamericana, (UNITEC), Honduras, *aliciareyes@unitec.edu**, *silviolazaro@unitec.edu*, *aaronortiz2000@unitec.edu*

Abstract*– This study developed a convolutional neural network for the automated detection of dental conditions—specifically caries, restorations, root canals, and prostheses—on radiographic images using the Roboflow platform and deployed on the NVIDIA Jetson Nano. A dataset of 1,933 dental X-ray images, provided by the Bright Smile clinic and expanded to 4,639 through data augmentation techniques, was used. The model was trained and tested in five iterative versions, gradually incorporating each dental condition class. Preprocessing techniques, including a 20% increase in image saturation and targeted augmentations, significantly improved detection performance. The final model achieved a mean Average Precision (mAP) of 97.7%, with notable improvement in the identification of dental caries—previously the most challenging class. These results demonstrate that optimized preprocessing combined with YOLO-based training in Roboflow and deployment on Jetson Nano constitutes an effective pipeline for real-time dental diagnostics.

***Keywords*-- Dental radiographs, dental caries detection, convolutional neural networks, object detection, incremental training, image preprocessing, data augmentation, Roboflow, YOLO.**

I. Introduction

Artificial intelligence (AI) has become a transformative tool in medical imaging, enabling more accurate and efficient diagnostic processes across multiple healthcare domains. In dentistry, AI-driven image analysis has gained increasing attention due to its potential to support clinicians in the detection of common oral conditions, particularly dental caries. Early and accurate identification of carious lesions is essential to prevent disease progression, reduce treatment costs, and improve patient outcomes.

Dental caries remains one of the most prevalent oral diseases worldwide, affecting individuals across all age groups. Radiographic imaging is routinely used for diagnosis; however, the interpretation of dental radiographs can be challenging due to low contrast, subtle lesion appearance, and inter-observer variability. These challenges make caries detection an ideal candidate for automated object detection methods based on convolutional neural networks (CNNs), which are capable of learning complex visual patterns that may be difficult to perceive through conventional visual inspection.

Recent advances in deep learning and object detection frameworks, such as YOLO-based architectures, have demonstrated promising results in medical image analysis. Nevertheless, detecting dental caries remains more complex than identifying other dental structures such as restorations, prostheses, or root canal treatments. This difficulty arises from the radiolucent nature of carious lesions and their visual similarity to normal anatomical variations and image artifacts in grayscale radiographs.

In parallel, the growing availability of edge-computing platforms has enabled the deployment of AI models in portable and real-time diagnostic systems. Embedded devices with GPU acceleration provide sufficient computational capability to execute deep learning models efficiently while maintaining low power consumption, making them suitable for clinical and point-of-care applications.

This study proposes an automated dental condition detection pipeline based on a convolutional neural network trained using a YOLO-based framework and managed through the Roboflow platform. The model is designed to detect multiple dental conditions—including caries, restorations, root canals, and prostheses—from radiographic images and is deployed on an embedded GPU platform for real-time inference.

A key contribution of this work is the analysis of image preprocessing strategies, particularly image saturation enhancement and data augmentation, to improve detection performance for dental caries, which represents the most challenging class. The results demonstrate that optimized preprocessing combined with iterative model training significantly enhances detection accuracy, supporting the feasibility of AI-assisted dental diagnostics in real-world clinical environments.

II. Related Work

The application of artificial intelligence to dental imaging has expanded significantly in recent years, with multiple studies exploring machine learning and deep learning techniques to support automated diagnosis from radiographic images. Early research demonstrated that learning-based approaches can enhance diagnostic quality assessment and support the identification of dental caries in bitewing and periapical radiographs, highlighting the potential of artificial intelligence in dentistry [1].

More recent works have focused on the direct application of deep learning models, particularly convolutional neural networks (CNNs), for automated analysis of radiographic images. Several IEEE-published studies report promising results in detecting complex visual patterns using CNN-based architectures and modern object detection frameworks, achieving high accuracy under controlled experimental conditions [2]–[4]. In this context, platforms such as Roboflow have been employed to support dataset management, annotation, augmentation, and comparative evaluation of YOLO-based models, demonstrating their usefulness in applied computer vision pipelines [3]. These studies confirm the suitability of deep learning for image-based detection tasks; however, they also indicate that the reliable identification of subtle features—such as early-stage dental caries—remains challenging due to low contrast and visual similarity with surrounding structures.

Object detection frameworks such as YOLO and related architectures have benefited from large-scale benchmarks and datasets originally developed for general-purpose object detection, such as the Microsoft COCO dataset. These benchmarks played a key role in advancing detection accuracy and robustness across domains and later enabled successful transfer of object detection models to medical imaging applications [5].

Beyond dentistry-specific applications, comprehensive surveys on deep learning–based object detection emphasize that model performance is strongly influenced by data preprocessing, augmentation strategies, and annotation quality rather than architecture alone [6], [7]. These findings highlight the importance of optimizing the entire training pipeline, particularly when detecting small or visually subtle objects in complex imaging scenarios.

The feasibility of deploying deep learning models on embedded and edge-computing platforms has also been investigated. Comparative studies evaluating low-power embedded devices, including Raspberry Pi–based systems, report limitations in processing capability and real-time performance for deep learning inference [8]. These results motivate the use of more capable embedded platforms when real-time object detection is required in applied computer vision systems.

Despite the advances reported in the literature, several gaps remain. Many existing studies prioritize overall accuracy metrics without systematically analyzing the impact of preprocessing strategies on the detection of visually subtle targets, such as dental caries in radiographic images. Furthermore, the deployment of multi-class detection models optimized for real-time inference in practical environments is often not addressed. These limitations motivate the present study, which focuses on optimized preprocessing, incremental training, and embedded deployment for automated dental condition detection.

III. METHODOLOGY

This section describes the methodology adopted to develop and evaluate the proposed automated dental condition detection system. The approach is structured into four main stages: dataset preparation and annotation, image preprocessing and augmentation, incremental training strategy, and embedded deployment for real-time inference.

A. Dataset Preparation and Annotation

The dataset used in this study consists of dental radiographic images containing multiple dental conditions, including restorations, root canal treatments, prostheses, and dental caries. Images were collected from a clinical environment and manually annotated using bounding boxes to identify the spatial location of each target condition.

Annotation was performed using the Roboflow platform, which provides tools for dataset management, labeling, and version control. The use of Roboflow enabled consistent annotation workflows and facilitated iterative refinement of the dataset through augmentation and retraining cycles.

To ensure consistency during the labeling process, standard clinical references were considered as guidance for lesion identification. In particular, the classification framework proposed by Black was used as a reference to support consistent interpretation of carious lesion location and severity across different tooth surfaces.

B. Image Preprocessing and Data Augmentation

Dental radiographs typically present low contrast and subtle intensity variations, especially in the case of early-stage carious lesions. To address these challenges, an image preprocessing stage was incorporated prior to model training.

A key preprocessing strategy evaluated in this work was saturation enhancement. Radiographic images were modified by increasing saturation levels by 20%, which improved the visibility of radiolucent features associated with dental caries. This enhancement facilitated clearer differentiation between healthy tooth structures and carious regions during both annotation and training.

In addition to saturation adjustment, data augmentation techniques were applied to increase dataset diversity and improve model generalization. The augmentation pipeline included horizontal and vertical flipping, brightness and exposure variations, random cropping, and automatic orientation. All images were resized to a fixed resolution of 640×640 pixels to meet the input requirements of the object detection model.

The relevance of preprocessing and controlled variation in image conditions is supported by prior CNN-based vision studies that demonstrated improved detection performance under heterogeneous capture scenarios [9].

C. Incremental Training Strategy

An incremental training strategy was adopted to progressively increase model complexity and analyze detection performance as new classes were introduced. Rather than training a multi-class model from the outset, the network was trained through successive stages, each incorporating an additional dental condition.

The first training stage focused exclusively on detecting dental restorations. Subsequent stages incrementally introduced root canal treatments, prostheses, and finally dental caries. This strategy enabled controlled evaluation of model behavior and helped identify which classes posed greater detection challenges.

The incremental approach aligns with prior work on staged development and refinement of complex systems, where progressive integration of components improves stability and performance [10]. In the context of object detection, this approach also facilitated analysis of class interference and model generalization as the number of detectable conditions increased as shown in Fig. 1.

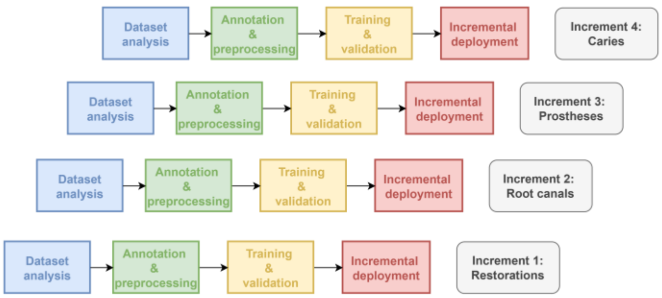

Fig. 1 Incremental training workflow adopted in this study, where dental conditions are progressively incorporated across successive training stages.

D. *Model Architecture and Training Configuration*

The object detection model was developed using a YOLO-based convolutional neural network architecture. This architecture was selected due to its balance between detection accuracy and computational efficiency, which is essential for real-time applications.

Training was conducted using Roboflow’s cloud-based infrastructure, enabling systematic experimentation with preprocessing configurations, augmentation strategies, and hyperparameters. The dataset was split into training, validation, and testing subsets following a 70%–20%–10% ratio. Model performance was evaluated using the mean Average Precision (mAP) metric, which is widely adopted for assessing object detection accuracy [11].

E. *Embedded Deployment for Real-Time Inference*

To evaluate the feasibility of real-time dental condition detection, the final trained model was deployed on an embedded GPU platform. The system was implemented using an NVIDIA Jetson Nano development kit, which provides GPU acceleration suitable for executing deep learning inference under low power consumption constraints [12].

This deployment stage focused on verifying that the optimized model could maintain reliable detection performance within the computational limitations of embedded hardware. The results demonstrate that the proposed pipeline is not limited to offline analysis and can be integrated into practical, real-time diagnostic environments.

F. *Ethical Considerations*

All radiographic images used in this study were anonymized prior to annotation and model training. No personally identifiable patient information was included in the dataset. The images were used exclusively for research purposes.

IV. Results and Discussion

This section presents and discusses the experimental results obtained from the proposed incremental training strategy for automated dental condition detection. Results are analyzed across successive training increments to evaluate models.

A. *Restoration Detection (Increment 1)*

In the first training increment, a convolutional neural network was trained exclusively for the automatic detection of dental restorations in radiographic images. This initial stage aimed to establish a baseline model capable of identifying well-defined and high-contrast structures commonly present in dental radiographs.

As illustrated in Fig. 2, a total of approximately 807 restorations were manually annotated and used for model training. The dataset was partitioned into training, validation, and testing subsets using a 70%–20%–10% split, respectively. This data distribution was adopted as a default configuration and consistently applied across all subsequent training increments to ensure comparability between models.

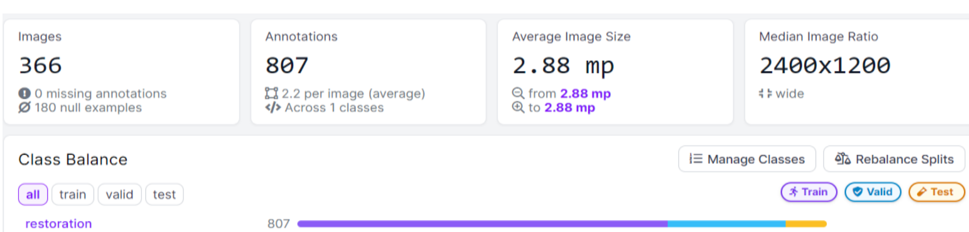

Fig. 2 Annotated dental radiographic images used in the first training increment, showing bounding boxes corresponding to dental restorations.

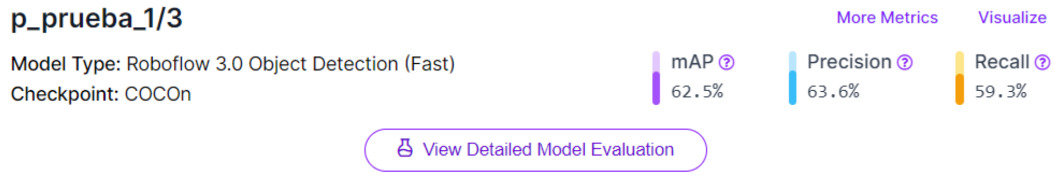

Fig. 3 Detection performance of the convolutional neural network trained for restoration detection, evaluated after the first incremental training stage.

At the end of this first training stage, which focused solely on the detection of restorations (caps) in dental radiographs, the model achieved a mean Average Precision (mAP) of 62.5%, as shown in Fig. 3. These results indicate that the network successfully learned representative visual features associated with restorations, providing a stable foundation for the incremental incorporation of additional dental conditions in later stages.

The first version of the neural network was evaluated to qualitatively analyze its detection behavior. As shown in Fig. 4, although the model can identify dental restorations, some orthodontic brackets are incorrectly classified as restorations during this initial training stage.

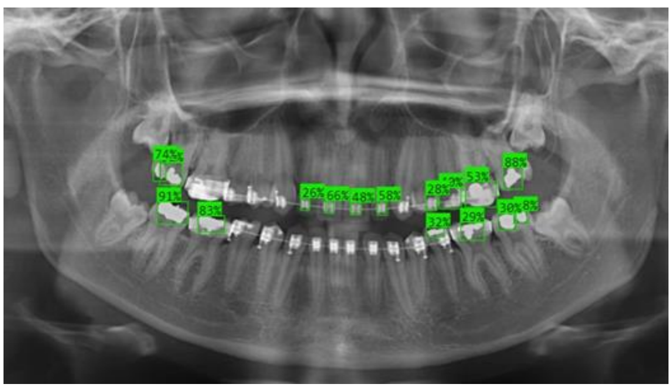

Fig. 4 Example of detection results obtained after the first training increment, illustrating correct restoration detection and misclassification of orthodontic brackets as restorations.

B. Root Canal Detection (Increment 2)

In the second training increment, the convolutional neural network was extended to incorporate root canal treatments as an additional detection class, while preserving the previously learned restoration class. This stage aimed to evaluate the model's ability to generalize to a new dental condition and to analyze the impact of multi-class learning on overall detection performance.

A total of approximately 702 root canal treatments were manually annotated and added to the dataset used in the first training increment. These new annotations were combined with the existing restoration labels, allowing the model to learn both classes simultaneously under the same training configuration.

Upon completion of this training stage, the model demonstrated an improvement in detection performance compared to the first increment. The introduction of the root canal class contributed to a more robust feature representation, resulting in an increase in the mean Average Precision (mAP) to 70.1%, as illustrated in Fig. 5. This improvement indicates that the model successfully adapted to the additional class without degrading its ability to detect restorations.

Qualitative results further confirm the successful incorporation of the root canal class. As shown in Fig. 6, the neural network is capable of identifying root canal treatments despite their limited visual contrast in some radiographic images.

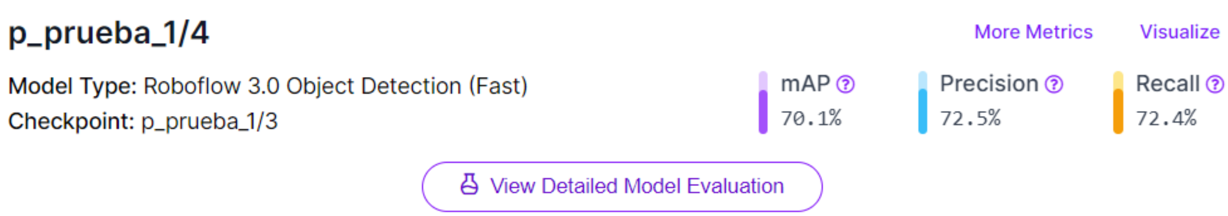

Fig. 5 Detection performance of the neural network after the second training increment, showing an improvement in mean Average Precision (mAP) following the incorporation of the root canal class.

These results demonstrate that the incremental training strategy enables the model to progressively learn new dental conditions while maintaining stable performance across previously learned classes.

Based on these findings, the next training stage was prepared to incorporate dental prostheses as an additional class, allowing further evaluation of model scalability and performance evolution.

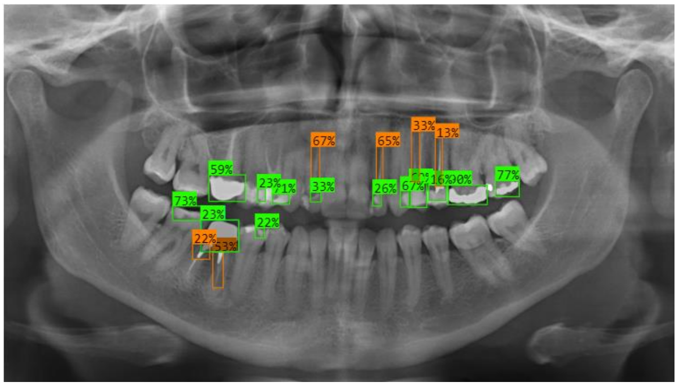

Fig. 6 Example of detection results obtained after the second training increment, illustrating successful identification of root canal treatments alongside dental restorations.

C. *Prostheses Detection (Increment 3)*

In the third training increment, dental prostheses were incorporated as an additional detection class, together with the previously learned restorations and root canal treatments. This stage aimed to evaluate the model's ability to handle increased class diversity and to analyze the impact of additional visual complexity on overall detection performance.

Approximately 529 dental prostheses were manually annotated and integrated into the existing dataset. The distribution of labeled instances for this third version of the model is illustrated in Fig. 7.

After completing the training process, the model exhibited a moderate improvement in performance compared to the previous increment. As shown in Fig. 8, the mean Average Precision (mAP) reached 73.7%, representing a smaller increase relative to the gains observed in earlier stages.

Qualitative detection results further illustrate the behavior of the model after incorporating dental prostheses. As shown in Fig. 9, the neural network is capable of simultaneously detecting restorations and prostheses in panoramic radiographs, correctly differentiating between both classes without significant confusion.

Despite the moderate improvement in mAP, the qualitative results demonstrate that the model successfully distinguishes prosthetic structures from restorations, even when both present similar radiographic appearances.

The limited improvement in mAP can be attributed to the relatively small number of prosthesis instances available for training. The reduced representation of this class may have constrained the network's ability to learn robust and discriminative features for prostheses while simultaneously maintaining accurate detection of restorations and root canal treatments. Additionally, the introduction of a third class increased classification complexity, requiring the model to distinguish between dental conditions with partially overlapping visual characteristics.

Despite the modest increase in mAP, qualitative evaluation indicates that the model is capable of correctly differentiating between prostheses and restorations, without significant class confusion. These results suggest that the incremental training strategy remains effective, and that further performance improvements can be expected in subsequent stages as additional data and more challenging classes are incorporated.

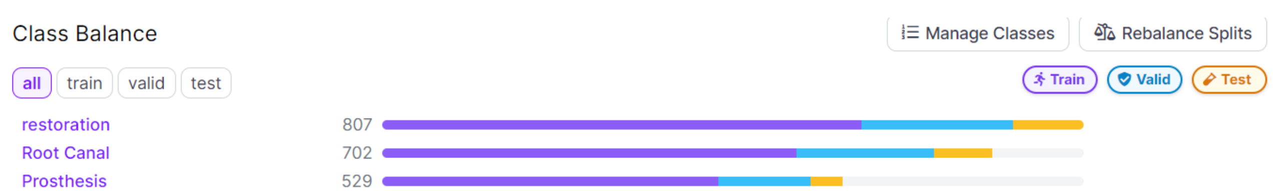

Fig. 7 Number of labeled instances used in the third training increment, including restorations, root canal treatments, and dental prostheses.

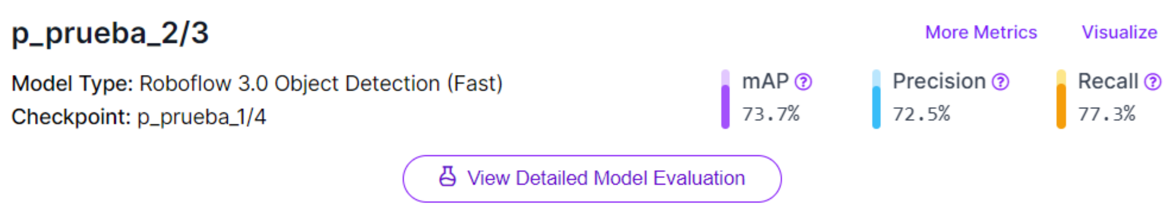

Fig. 8 Detection performance of the neural network after the third training increment, showing the mean Average Precision (mAP) achieved following the incorporation of dental prostheses.

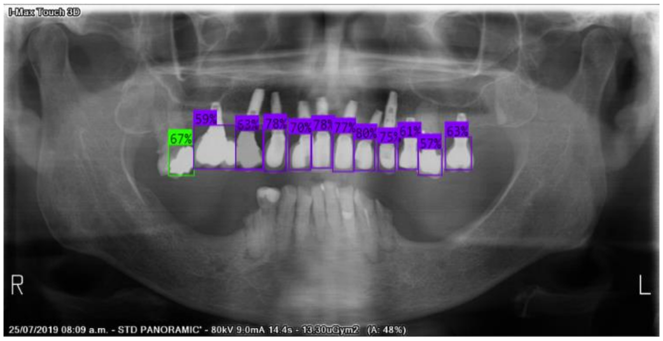

Fig. 9 Example of restoration and prosthesis detection after the third incremental training stage.

D. Caries Detection (Increment 4)

In the fourth and final training increment, dental caries were incorporated as an additional detection class, together with restorations, root canal treatments, and dental prostheses. This stage represents the most challenging scenario, as carious lesions typically exhibit low contrast, diffuse boundaries, and radiolucent characteristics that are difficult to distinguish in panoramic radiographs.

Approximately 781 carious lesions were manually annotated and added to the dataset. In addition, supplementary instances of prostheses and root canal treatments were included to improve class balance and reduce potential bias during training. The resulting distribution of labeled instances for this fourth version of the model is illustrated in Fig. 10.

Despite efforts to balance the dataset across all classes, the model exhibited a significant decrease in detection performance after incorporating dental caries. As shown in Fig. 11, the mean Average Precision (mAP) dropped considerably compared to previous training increments. This outcome indicates that dental caries pose a substantially greater detection challenge than other dental conditions within the proposed object detection framework.

Qualitative analysis confirms that the model struggles to reliably identify carious lesions in radiographic images. While large and well-defined cavities are occasionally detected, small or early-stage caries are frequently missed, as illustrated in Fig. 12. This behavior aligns with clinical observations, as the visual identification of caries in panoramic radiographs is also challenging for human experts.

Fig. 10 Distribution of annotated instances used in the fourth training increment, including restorations, root canal treatments, prostheses, and dental caries.

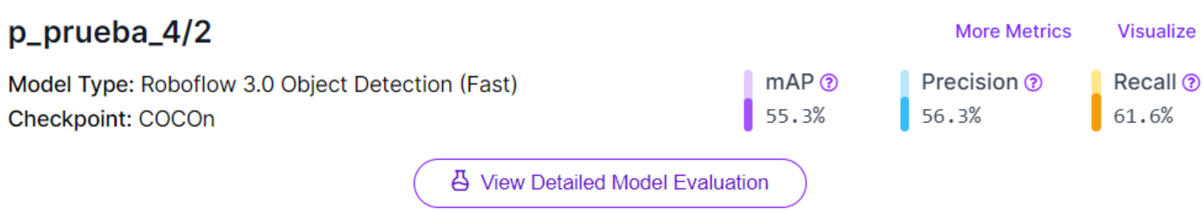

Fig. 11 Mean Average Precision (mAP) obtained after the fourth incremental training stage, illustrating the performance degradation following the incorporation of dental caries.

Additional experiments were conducted to improve caries detection by relabeling samples and modifying the distribution of training, validation, and testing datasets. However, these adjustments did not result in a performance improvement beyond an mAP of approximately 63%, reinforcing the conclusion that the current preprocessing and object detection configuration is insufficient for reliable caries detection.

These findings highlight the critical importance of advanced image preprocessing and enhancement techniques when addressing low-contrast dental conditions. Improving contrast, texture representation, and lesion visibility during preprocessing is essential to support both accurate annotation and effective model training, particularly for dental caries, which constitute the primary focus of this study.

The limited improvement in mAP can be attributed to the relatively small number of prosthesis instances available for training. The reduced representation of this class may have constrained the network's ability to learn robust and discriminative features for prostheses while simultaneously maintaining accurate detection of restorations and root canal treatments. Additionally, the introduction of a third class increased classification complexity, requiring the model to distinguish between dental conditions with partially overlapping visual characteristics.

Despite the modest increase in mAP, qualitative evaluation indicates that the model is capable of correctly differentiating between prostheses and restorations, without significant class confusion. These results suggest that the incremental training strategy remains effective, and that further performance improvements can be expected in subsequent stages as additional data and more challenging classes are incorporated.

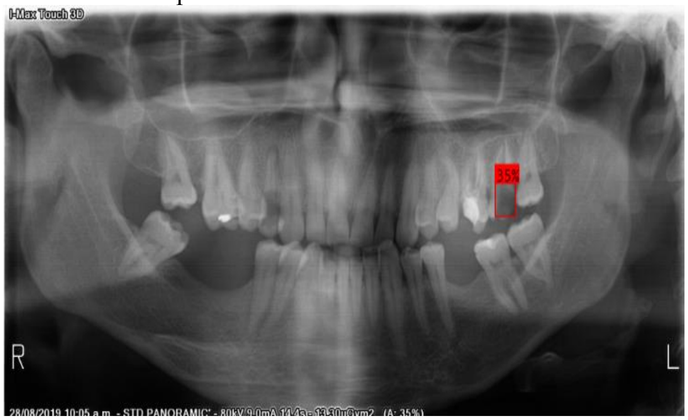

Fig. 12 Qualitative detection results for dental caries after the fourth training increment, showing successful detection of large cavities and failure to identify small or subtle lesions.

Optimizing Labeling and Object Detection Performance

Based on the results obtained in the fourth training increment, it was concluded that the primary limitation affecting caries detection was related to insufficient visual contrast in the radiographic images. Given that dental caries are inherently difficult to distinguish due to their radiolucent nature, improvements at the preprocessing stage were required to enhance feature visibility prior to labeling and training.

To address this issue, image saturation was increased by 20%, as illustrated in Fig. 13. This adjustment was applied to improve the visual differentiation of carious regions, facilitating more accurate annotation and more effective feature learning by the neural network.

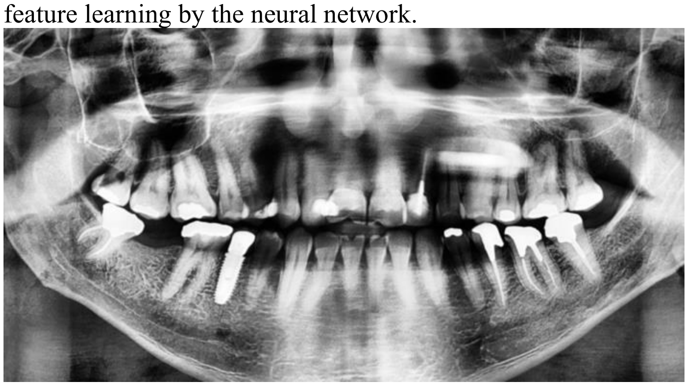

Fig. 13 Example of dental radiographs after applying a 20% saturation increase to enhance the visibility of carious lesions.

A new training version was subsequently developed using the same dataset as the fourth increment, but with all images preprocessed using the increased saturation level. The number of labeled instances remained unchanged, consisting of 809 restorations, 808 root canal treatments, 781 dental caries, and 759 prostheses, ensuring that performance differences could be attributed exclusively to preprocessing improvements rather than data volume.

Following this adjustment, a significant improvement in detection performance was observed. As shown in Fig.15, the mean Average Precision (mAP) increased to 89.9%, demonstrating that saturation enhancement substantially improved the model's ability to learn discriminative features, particularly for dental caries.

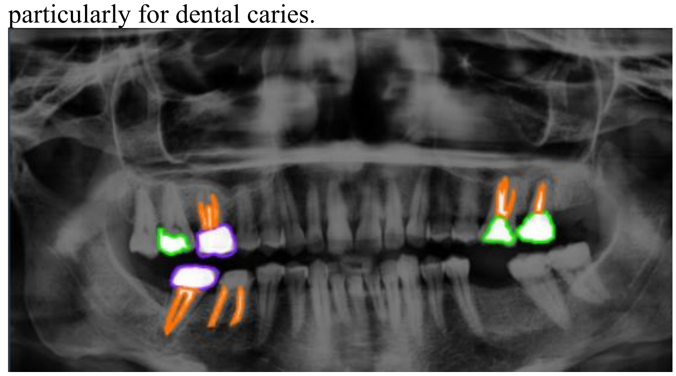

Fig. 14 Preprocessing and data augmentation strategies applied during the final training stage to improve model generalization.

To further optimize performance, additional data augmentation techniques were incorporated, as summarized in Fig. 14. These augmentations were designed to increase dataset variability and improve model generalization. The applied preprocessing and augmentation strategies included:

Preprocessing

- Automatic orientation correction
- Resize with fit and black padding to 640 × 640 pixels

Data Augmentation

- Horizontal and vertical flipping
- Random cropping with 0% minimum zoom and 20% maximum zoom
- Brightness variation between –15% and +15%
- Exposure variation between –10% and +10%

After applying these augmentations and retraining the network, the model achieved a final mAP of 97.7%, as presented in Fig. 16. This result represents the highest performance obtained throughout the incremental training process and indicates a substantial improvement in caries detection accuracy alongside the other dental conditions.

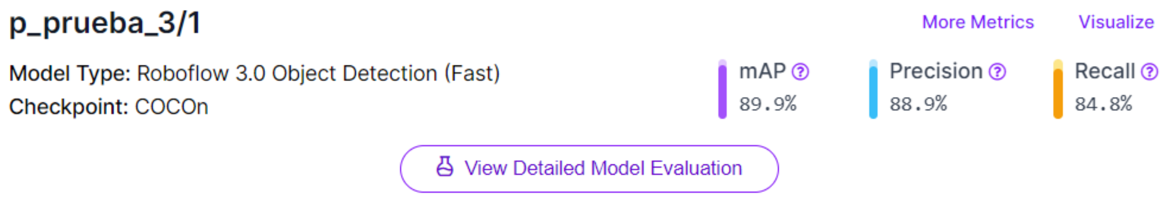

Fig. 15 Mean Average Precision (mAP) obtained after retraining the model using saturation-enhanced images.

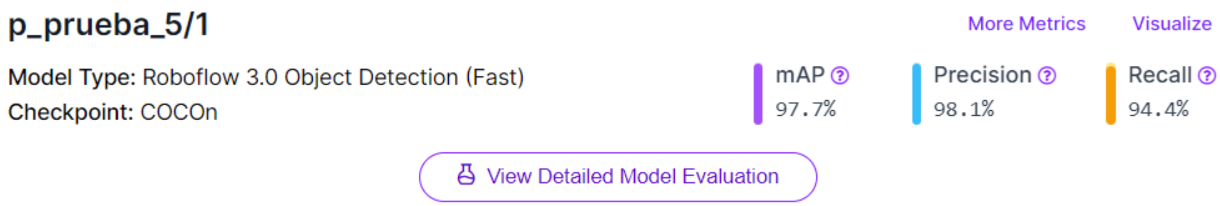

Fig. 16 Final detection performance achieved after applying saturation enhancement and data augmentation, reaching an mAP of 97.7%.

E. Limitations

Despite the promising detection performance achieved in this study, several limitations should be acknowledged. First, the dataset was obtained from a single clinical source, which may limit the diversity of radiographic acquisition conditions and patient characteristics. Second, the study focused on panoramic radiographs only; therefore, model performance on bitewing or periapical radiographs remains unknown.

Additionally, although the proposed preprocessing strategy significantly improved detection performance, further validation using independent external datasets and cross-validation strategies is required to fully assess model robustness and generalizability.

Finally, small and early-stage carious lesions continue to represent the most challenging detection scenario due to their subtle radiographic appearance.

Qualitative results further support these findings. As illustrated in Fig. 17, saturation enhancement combined with data augmentation enabled more effective class separation and improved detection consistency, particularly for dental caries. By enhancing contrast and highlighting subtle radiographic features, the preprocessing strategy facilitated more accurate object localization and classification, ultimately improving the overall robustness of the detection system.

These results confirm that image preprocessing and augmentation play a critical role in the successful application of object detection models to dental radiographic analysis, especially for visually subtle conditions such as caries.

The final detection performance achieved in this study is comparable to or exceeds several previously reported CNN-based dental detection approaches described in the literature. However, direct comparison remains difficult due to differences in datasets, annotation methodologies, radiographic modalities, and evaluation metrics across studies.

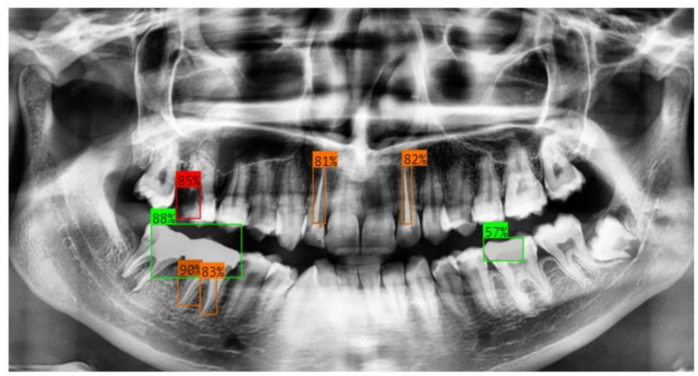

Fig. 17 Qualitative detection results illustrating improved class separation and caries detection after saturation enhancement and data augmentation.

V. CONCLUSIONS

In this research, the RoboFlow model demonstrated that detecting restorations, prostheses, and root canals did not present as much difficulty compared to detecting cavities. This difference is due to several factors related to the nature of the images and the ability of the neural network to interpret variations in gray scale. Cavities present a greater challenge because the affected areas can be subtle and are often confused with other tone variations in the image, especially in grayscale examples.

During the training process, numerous adjustments to the parameters of the image distributions were attempted for the training, validation, and testing data sets. Despite these efforts, the different versions developed in RoboFlow failed to exceed a performance of 63% compared to the first versions of the neural networks. These initial versions did not yet label the images in a way that the RoboFlow model could effectively use them as examples.

A significant advance was achieved by saturating the images by 20%. This adjustment in saturation allowed the model to detect the classes more easily and especially improved the detection of cavities. Cavities, which were previously difficult to identify, were only detected when they were quite large. Additionally, the network sometimes confuses tooth shadows with cavities, reducing the model's accuracy. Saturation improved the contrast of the images, making cavities more visible and distinguishable from other dental structures.

In addition to adjusting saturation, augmentation techniques were implemented in the latest version of the model. These techniques included variations in brightness, contrast, and other transformations that helped diversify the training data set. As a result, an expanded data set was generated that went from 1,933 original images, provided by the Bright Smile clinic, to a total of 4,639 images after preprocessing and augmentations. This significant increase in the amount of data allowed the neural network to train with a greater variety of examples, improving its ability to generalize and detect objects in new images.

Thanks to these improvements, the model achieved a very high mAP, which indicates high precision in object detection. The combination of image saturation and augmentation techniques proved to be crucial in overcoming initial limitations and significantly improving model performance.

It is strongly recommended that when using RoboFlow in future research involving radiography, these strategies be considered. Saturating images can improve the visibility of radiolucent objects, which appear as darker shades of gray on x-rays. In addition, it is essential to carefully label the entire contour of the objects that you want to train in the neural network. This approach will ensure that the network fully understands the shape and characteristics of objects, which will improve its detection ability and reduce the possibility of confusion. Implementing these practices will allow for the development of more precise and efficient detection models in the field of dental radiography.

References

[1] M. A. Barayan, A. A. Qawas, A. S. Alghamdi, T. S. Alkhallagi, R. A. Al-Dabbagh, G. A. Aldabbagh, and A. I. Linjawi, “Effectiveness of machine learning in assessing the diagnostic quality of bitewing radiographs,” Applied Sciences, vol. 12, no. 19, p. 9588, 2022, doi: 10.3390/app12199588.

[2] J. L. Ordoñez-Avila, E. R. T. Maldonado, and I. Magomedov, “Water generation based on condensation controlled by gray scale and artificial vision,” in Proc. 2022 Int. Conf. Information, Control, and Communication Technologies (ICCT), Astrakhan, Russia, 2022, pp. 1–5, doi: 10.1109/ICCT56057.2022.9976774.

[3] J. S. Aleman and A. M. Reyes-Duke, “Performance analysis of Roboflow 3.0 and YOLO-NAS in PCB defect detection compared to YOLOv9 and RT-DETR using augmented image dataset,” in Proc. 10th Int. Conf. Control and Robotics Engineering (ICCRE), Nagoya, Japan, 2025, pp. 229–234, doi: 10.1109/ICCRE65455.2025.11093496.

[4] C. A. Bardales Cortés and A. M. Reyes-Duke, “A comparative study of CNN trainings for the detection of E. coli, P. aeruginosa, and S. aureus: A Raspberry Pi–based prototype,” in Proc. IEEE 41st Central America and Panama Convention (CONCAPAN XLI), Tegucigalpa, Honduras, 2023, pp. 1–6, doi: 10.1109/CONCAPANXLI59599.2023.10517539.

[5] T.-Y. Lin et al., “Microsoft COCO: Common objects in context,” in Proc. Eur. Conf. Comput. Vis. (ECCV), Zurich, Switzerland, 2014, pp. 740–755, doi: 10.1007/978-3-319-10602-1_48.

[6] K. Sharma, S. Rawat, D. Parashar, S. Sharma, S. Roy, and S. Sahoo, “State-of-the-art analysis of multiple object detection techniques using deep learning,” Int. J. Adv. Comput. Sci. Appl., vol. 14, pp. 527–534, 2023, doi: 10.14569/IJACSA.2023.0140657.

[7] H. M. Ahmad and A. Rahimi, “Deep learning methods for object detection in smart manufacturing: A survey,” Journal of Manufacturing Systems, vol. 64, pp. 181–196, 2022, doi: 10.1016/j.jmsy.2022.06.011.

[8] G. Dewantoro, J. Mansuri, and F. D. Setiaji, “Comparative study of computer vision-based line followers using Raspberry Pi and Jetson Nano,” Jurnal Rekayasa Elektrika, vol. 17, no. 4, 2021, doi: 10.17529/jre.v17i4.21324.

[9] J. A. Menjívar Rivera and A. M. Reyes-Duke, “Convolutional neural network for the detection of cocoa maturity with an approach for the analysis of images captured at different distances,” in Proc. Int. Conf. Science, Technology and Innovation (CONICIETI), ser. Engineering Headway, vol. 12, Trans Tech Publications Ltd., Oct. 2024, pp. 3–8, doi: 10.4028/p-JM3tGd.

[10] B. Boehm and J. A. Lane, “Using the incremental commitment model to integrate system acquisition, systems engineering, and software engineering,” *CrossTalk*, vol. 19, no. 10, pp. 4–9, 2007.

[11] S. de Witte, O. Strafforello, and J. van Gemert, “Do object detection localization errors affect human performance and trust?” *arXiv preprint arXiv:2401.17821*, 2024.

[12] NVIDIA, “Jetson Nano developer kit,” NVIDIA Developer Documentation, 2024. [Online]. Available: https://developer.nvidia.com/embedded/jetson-nano