

Effectiveness of lightweight convolutional neural networks for detecting the relationship between the mandibular third molar and the inferior alveolar canal on panoramic radiographs

Ali Afzoon Khiyavi¹, Abolfazl Shiri Varnamkhasti^{2,3,*}, Maryam Tofangchiha⁴, Ali Labafchi^{3,5}

Abstract

Objective: This study aimed to develop and evaluate lightweight convolutional neural networks (CNNs) capable of automatically localizing the mandibular third molar (M3) and classifying its relationship with the inferior alveolar canal (IAC) on panoramic radiographs.

Methods: A total of 609 panoramic radiographs (containing 899 M3s) were analyzed in two stages. First, 82 panoramic images (134 M3s) were used to fine-tune a pre-trained EfficientDet model for automatic M3 localization. The detected regions were standardized to include the IAC and preprocessed through resizing, contrast enhancement, and mirroring. For ground-truth labeling, the presence or absence of M3–IAC contact was determined by an experienced oral and maxillofacial radiologist based on established panoramic radiographic criteria. Second, a custom lightweight CNN was trained on 527 panoramic radiographs (765 M3s) to classify M3–IAC contact (contact = 1, no contact = 0). Model performance was compared with a pre-trained ResNet50 architecture using accuracy, sensitivity, specificity, precision, and F1 score.

Results: The detection model achieved 100% accuracy with an intersection-over-union (IoU) of 87.9%. Compared to the ResNet50 benchmark model, the lightweight CNN demonstrated comparable overall accuracy (87.5%). However, the lightweight CNN outperformed ResNet50 in specificity (90.4% versus 86.9%) and precision (93.4% versus 88.7%), while ResNet50 exhibited a slightly higher mean sensitivity (88.3% versus 86.2%).

Conclusions: Lightweight CNNs can achieve diagnostic performance comparable to large pre-trained networks while requiring less training time and computational power. The proposed model enables automated, efficient, and clinically feasible detection of the M3–IAC relationship on panoramic radiographs.

Keywords: Artificial intelligence, Computer-assisted image processing, Convolutional neural networks, Inferior alveolar canal, Mandibular third molar, Panoramic radiographs

Introduction

Extraction of mandibular third molars (M3s) is one of the most common oral surgical procedures performed by dentists and oral and maxillofacial surgeons (1). Despite being routine, this procedure is associated with

several potential complications, including postoperative pain and swelling, alveolar osteitis, infection, paresthesia of the tongue or lip, and even mandibular fracture (2-5). Accurate preoperative assessment of the anatomical relationship between the M3 roots and the inferior alveolar canal (IAC) is essential for minimizing the risk of nerve injury and for proper surgical planning (6, 7). Therefore, improving the diagnostic accuracy of the M3-IAC relationship on panoramic radiographs remains clinically important.

The identification of the IAC and its relationship with the M3 roots requires a high level of expertise in radiographic interpretation. However, the limited availability of trained oral and maxillofacial radiologists, combined with interobserver variability among clinicians, can affect the reliability of diagnoses (8, 9). To overcome these challenges, artificial intelligence (AI) has emerged as a promising tool in dental imaging,

¹ Department of Electrical Engineering, Iran University of Science and Technology, Tehran, Iran.

² Oral and Dental Diseases Research Center, Kerman University of Medical Sciences, Kerman, Iran.

³ Student Research Committee, Faculty of Dentistry, Kerman University of Medical Sciences, Kerman, Iran.

⁴ Department of Oral and Maxillofacial Radiology, Dental Caries Prevention Research Center, Qazvin University of Medical Sciences, Qazvin, Iran.

⁵ Department of Oral and Maxillofacial Surgery, Faculty of Dentistry, Kerman University of Medical Sciences, Kerman, Iran.

*Corresponding Author: Abolfazl Shiri Varnamkhasti
Email: a.shiri@kmu.ac.ir

Accepted: 7 March 2025. Submitted: 22 November 2025.



demonstrating excellent performance in image recognition and classification tasks (10, 11).

Among AI techniques, convolutional neural networks (CNNs) have demonstrated remarkable success in medical image analysis, including dental radiography (12). However, most existing models rely on manually selected regions of interest (ROIs) and use large, computationally intensive architectures (e.g., VGG16, ResNet), which limit their practicality in clinical settings due to high hardware demands and slow inference times (13, 14). In contrast, lightweight or compact CNN architectures can offer faster and more efficient alternatives without compromising accuracy.

The present study aimed to develop and evaluate small, efficient CNN models capable of automatically localizing mandibular third molars on panoramic radiographs and classifying their relationship with the IAC.

Materials and methods

All procedures were conducted in accordance with relevant institutional and international ethical standards. The study protocol was approved by the ethics committee of Qazvin University of Medical Sciences (IR.QUMS.REC.1400.212). All panoramic radiographs were anonymized, and personal identifiers, including name, age, and gender, were removed before analysis.

Study design and workflow overview

This study consisted of two main phases: (1) automatic localization of mandibular third molars (M3s) using a detection model based on fine-tuning the EfficientDet architecture, and (2) classification of the anatomical relationship between the M3 and the inferior alveolar

canal (IAC) using custom-developed CNNs that were trained and validated for this purpose.

A summary of the workflow is illustrated in Figure 1. An experienced oral and maxillofacial radiologist established the gold standard for this study by annotating all panoramic radiographs. The radiologist drew bounding boxes around each mandibular third molar (M3) and determined the presence or absence of contact between the M3 and the inferior alveolar canal (IAC). These annotated images were then used to fine-tune an EfficientDet-based object detection model for automatic localization of the M3. The automatically detected regions were subsequently expanded to ensure inclusion of the IAC, preprocessed, and used as input for a custom lightweight CNN. This CNN was trained to classify the M3–IAC relationship as either contact (label = 1) or non-contact (label = 0).

Dataset characteristics

Panoramic radiographs were retrospectively collected from two sources between April 15, 2021, and September 24, 2021. The first dataset was obtained from the Department of Oral and Maxillofacial Radiology at Qazvin Dental School and was acquired using a Rayscan α -P system (Ray, Seoul, South Korea; 73 kVp, 10 mA, 14 s). The second dataset was collected from a private dental radiology clinic and was acquired using a Cranex 3D unit (Soredex, Helsinki, Finland; 70 kVp, 16 mA, 15 s).

The dataset consisted of 609 panoramic radiographs, which contained a total of 899 mandibular third molars (M3s). All patients were ≥ 20 years and had fully developed M3 roots. The exclusion criteria included immature M3 buds, horizontal angulation $\geq 90^\circ$, evidence of buccolingual angulation, and abnormal tooth morphology. Teeth with previous endodontic




Figure 1. Workflow of the automated M3–IAC contact detection system. The model produced a continuous output between 0 and 1, corresponding to non-contact (label 0) and contact (label 1), respectively

treatment or orthodontic brackets were not excluded from the study.

To minimize device-related variability, all panoramic images acquired with the Cranex 3D unit (in RGB format) were converted to grayscale to match the output format of the Rayscan α -P® system.

The dataset was divided into two parts corresponding to the detection and classification phases of the study. For the detection phase, 82 panoramic radiographs encompassing 134 half-panoramic regions were used to train and fine-tune the object detection model for automatic localization of mandibular third molars (M3s). For the classification phase, 527 panoramic radiographs comprising 765 half-panoramic regions were included, with 396 labeled as contact and 369 as non-contact between the M3 and the inferior alveolar canal.

All radiographs were independently annotated by a certified oral and maxillofacial radiologist with over 25 years of experience, who served as the gold standard. For the detection dataset, the radiologist manually drew bounding boxes around each mandibular third molar (M3). For the classification dataset, the radiologist labeled each M3 based on the presence of one or more of the seven established panoramic radiographic signs indicating proximity or contact with the inferior alveolar canal (IAC) including; darkening of the root, deflection of the root, narrowing of the IAC, a dark and bifid apex, an interruption in the cortical border of the IAC, diversion of the IAC and narrowing of the root (15).

Detection model

To automate the localization of mandibular third molars (M3), an EfficientDet-D1 model (16) was fine-tuned using the TensorFlow Object Detection API. The base model, initially trained on the COCO 2017 dataset (17), was adapted through transfer learning to detect the terminal mandibular tooth within each half-panoramic region.

The model generated four normalized coordinates (Ymin, Xmin, Ymax, Xmax) defining the predicted bounding box. These normalized values allowed automatic cropping of the M3 region regardless of image dimensions, ensuring compatibility across different input sizes and scanner outputs. The cropped regions were then used as inputs for the classification phase.

Training was conducted as a single-class detection task with a batch size of 1, a learning rate of 0.001, and 20,000 training steps. Localization performance was assessed using the mean intersection-over-union (IoU), measuring the spatial overlap between the model-

predicted bounding boxes and those manually annotated by the radiologist.

Image preprocessing

A standardized image preprocessing sequence (pipeline) was applied to all extracted regions of interest (ROIs) before classification to ensure consistency and optimize model performance. This sequence consisted of automated ROI expansion, geometric normalization, and intensity standardization. First, the radiologist-defined or model-generated bounding boxes were systematically expanded in the inferior and posterior directions to ensure consistent inclusion of the inferior alveolar canal (IAC), while excluding the coronal portion of the mandibular third molar. This step ensured that all ROIs encompassed the anatomically relevant region for assessing M3–IAC proximity. The expanded ROIs were subsequently resized to 220 × 220 pixels, providing uniform input dimensions for the convolutional neural network.

Image quality was further improved by applying contrast adjustment (factor = 1.5) using the PILLOW library, which enhanced the visibility of root contours and canal borders. To ensure consistent anatomical orientation, all right-side third molars were horizontally flipped, enabling the model to learn direction-independent features and reducing variability within the dataset.

During training, random shuffling (the process of randomly rearranging the order of data in a dataset) and data augmentation techniques were used to reduce overfitting and make the model more adaptable to new data. These techniques included small rotations, flips, and changes in intensity. By applying these methods, the dataset became more balanced and standardized, reducing the need for manual adjustments during both the training process and the model's inference.

Classification model

A custom lightweight convolutional neural network (CNN) was designed to perform binary classification of the anatomical relationship between the mandibular third molar (M3) and the inferior alveolar canal (IAC). The network architecture comprised four sequential blocks of layers, with each block using more filters to analyze the image. Each block included techniques like batch normalization, ReLU activation, max pooling, and dropout (with a dropout rate of 0.25). These techniques helped the network learn better features from the data while also reducing the risk of overfitting.

To enhance feature extraction in the early stages, larger convolutional kernel sizes (7×7 and 5×5) were used in the first layers. This allowed the model to capture larger structural patterns, such as the canal's path and root shape, while minimizing the number of complex transformations. Zero-padding, a technique used in CNNs to preserve the spatial dimensions of an image after applying convolution operations, was applied to maintain the spatial resolution at the image borders. The output from the last convolutional block was flattened and passed into a fully connected layer with eight nodes (neurons). A sigmoid activation function was applied to the outputs from the neurons, compressing them into a binary result that indicates either a contact (1) or no contact (0) relationship.

Model training was performed using Python 3.9 and TensorFlow 2.5. The network was optimized with the Adam optimizer, using a learning rate of 0.001 and the default momentum parameters ($\beta_1 = 0.9$, $\beta_2 = 0.999$). The binary cross-entropy loss function was used for supervision. All experiments were conducted using an NVIDIA GTX 1650 GPU with 4 GB of VRAM.

Benchmarking lightweight CNNs: Comparison with ResNet50

To benchmark the performance of the proposed lightweight CNN, a state-of-the-art ResNet50 architecture (18) pre-trained on the ImageNet dataset was implemented as a comparison model. The pre-trained convolutional layers of ResNet50 were used as a deep feature extractor, with fully connected layers added afterwards to enable binary classification of the M3–IAC relationship. The ResNet50 model was trained on the same dataset used for the lightweight CNN and underwent identical preprocessing, augmentation, and evaluation procedures to ensure a fair comparison.

Due to the significantly higher computational and memory requirements of ResNet50 compared to the lightweight model, training was performed on a separate workstation equipped with an 8 GB GPU. All other training parameters, including optimization settings and loss functions, were kept consistent with those of the lightweight CNN to ensure that architectural differences remained the primary variable under comparison.

Training, validation, and evaluation of model performance

To get a more reliable estimate of model performance and reduce bias, each model was trained and evaluated

multiple times using random splits of the dataset. For each split, 80% was used for training, 5% for validation, and 15% for testing. During training, early stopping was applied based on validation loss to prevent overfitting and enhance the model's generalizability.

Model performance was evaluated using standard classification metrics, including accuracy, sensitivity, specificity, precision, and F1 score. The radiologist's annotations served as the gold standard for comparison.

The following metrics were calculated for each model:

Accuracy: It is the ratio of the number of correct predictions (both true positives and true negatives) to the total number of predictions.

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN}$$

Sensitivity: Sensitivity measures how well the model detects actual positives. It is the ratio of true positives to the total number of actual positives (true positives + false negatives).

$$Sensitivity = \frac{TP}{TP+FN}$$

Specificity: Specificity measures how well the model detects actual negatives. It is the ratio of true negatives to the total number of actual negatives (true negatives + false positives).

$$Specificity = \frac{TN}{TN+FP}$$

Precision: Precision measures how many of the positive predictions made by the model were actually correct. It is the ratio of true positives to the total predicted positives (true positives + false positives).

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

F1 Score: F1 Score is the harmonic mean of precision and sensitivity. It balances the trade-off between the two, especially when you need to consider both false positives and false negatives equally. A higher F1 score indicates better performance, especially in imbalanced datasets.

$$F1\ Score = \frac{TP}{TP + \frac{1}{2}(FP+FN)}$$

Receiver operating characteristic (ROC) curves and area under the curve (AUC) values were also used to assess the classifiers' ability to discriminate between classes at different thresholds. To reduce the effect of randomness in the training process, performance metrics were averaged over three separate training runs for both the lightweight CNN and the ResNet50 benchmark model, and the mean values were recorded.

Table 1. Performance metrics of the three lightweight CNN classification models

	Accuracy %	Sensitivity %	Specificity %	Precision %	F1 score %
Model 1	90.43	90.00	91.11	94.03	91.97
Model 2	86.96	88.06	85.41	89.39	88.72
Model 3	85.22	80.52	94.73	96.87	87.94
Mean	87.53	86.19	90.42	93.43	89.54

Results

Performance of the detection model

The fine-tuned EfficientDet-D1 model accurately localized all mandibular third molars (M3s) in the panoramic images, achieving a detection accuracy of 100% and a mean intersection-over-union (IoU) of 87.9%. Manual verification confirmed that all bounding boxes correctly included the inferior alveolar canal (IAC) within the region of interest (ROI). No false localizations or missed detections occurred in the test set.

Classification model performance

Three independently trained lightweight CNN models were evaluated on the test dataset, with their performance metrics summarized in Table 1. The top-performing model achieved an accuracy of 90.4%, sensitivity of 90.0%, specificity of 91.1%, precision of 94.0%, and an F1 score of 92.0%. The corresponding ROC curves showed AUC values ranging from 0.905 to 0.971 across the three models.

The mean values across the three runs were 87.5% for accuracy, 86.2% for sensitivity, 90.4% for specificity, 93.4% for precision, and 89.5% for F1 score. The performance results for the three models are shown in Figure 2.

Comparison with the benchmark ResNet50 model

To evaluate relative performance, a pre-trained ResNet50-based model was trained and tested under identical conditions. The results are summarized in Table 2 and Figure 3. The best ResNet50 model achieved an accuracy of 88.7%, sensitivity of 91.3%, specificity of 86.0%, precision of 86.9%, and an F1 score of 89.1%. The corresponding ROC curves yielded AUC values ranging from 0.921 to 0.964 across the three models.

The average values across ResNet50 models were 87.5% for accuracy, 88.3% for sensitivity, 86.9% for specificity, 88.7% for precision and 88.9% for F1 score.

Discussion

In this study, a fully automated sequence was developed to localize the mandibular third molar (M3) and classify its relationship to the inferior alveolar canal (IAC) on panoramic radiographs. The key finding was that a custom lightweight CNN achieved diagnostic performance comparable to a deeper, pre-trained ResNet50 model, while requiring significantly less memory (<1 MB), shorter training time (~3 minutes), and fewer computational resources. These advantages highlight its potential suitability for routine clinical




Figure 2. The performance results for the three lightweight CNN models.

Table 2. Performance metrics of the three ResNet50-based experimental models.

	Accuracy %	Sensitivity %	Specificity %	Precision %	F1 score %
Model 1	86.09	83.08	90.00	91.53	87.10
Model 2	88.70	91.30	85.96	86.89	89.08
Model 3	87.83	90.48	84.62	87.69	90.48
Mean	87.53	88.31	86.86	88.70	88.89

Predicted	Real Labels	
	TRUE	FALSE
	TRUE	54
FALSE	11	45

Predicted	Real Labels	
	TRUE	FALSE
	TRUE	53
FALSE	5	49

Predicted	Real Labels	
	TRUE	FALSE
	TRUE	57
FALSE	6	44

Figure 3. The performance results for the three ResNet50-based experimental models.

settings, where access to high-performance GPU hardware is often limited.

In the present study, several steps were taken to minimize bias and improve the reliability of the results. All radiographs were annotated by a single experienced oral and maxillofacial radiologist, providing a consistent gold standard for both detection and classification tasks. Sampling bias was addressed by using multiple randomized splits of the dataset. To prevent overfitting, dropout regularization, data augmentation, and early stopping based on validation performance were employed. Additionally, all bounding boxes generated automatically by the detection model were manually reviewed for accuracy before being used in the classification phase, further ensuring the quality and integrity of the training data.

The detection stage demonstrated excellent performance. The fine-tuned EfficientDet-D1 network successfully localized all M3s in the dataset, achieving 100% detection accuracy with a mean intersection-over-union (IoU) of 87.9%. Manual verification confirmed no missed teeth or incorrect bounding boxes. This high precision suggests that automated identification of the M3 region on panoramic radiographs is reliable, effectively eliminating the need for manual ROI selection, which is prone to operator variability (1). In contrast, previous studies often relied on manually cropped or defined ROIs (13), which limits reproducibility and scalability in clinical practice.

In the classification stage, the lightweight CNN models consistently demonstrated strong performance. The best-performing lightweight model achieved an accuracy of 90.4%, sensitivity of 90.0%, specificity of 91.1%, precision of 94.0%, and an F1 score of 92.0%. The mean performance across three independent runs remained comparably high, with accuracy at 87.5%, sensitivity at 86.2%, specificity at 90.4%, precision at 93.4%, and an F1 score at 89.5%. These findings indicate that a compact network can effectively differentiate between contact and non-contact cases using only panoramic radiographic features. The corresponding

ROC curves yielded AUC values ranging from 0.905 to 0.971, reflecting excellent discriminative capability.

Compared to the ResNet50 benchmark model, the lightweight CNN demonstrated comparable overall accuracy, with both achieving a mean accuracy of 87.5%. However, the lightweight CNN outperformed ResNet50 in specificity (90.4% versus 86.9%) and precision (93.4% versus 88.7%), while ResNet50 exhibited a slightly higher mean sensitivity (88.3% versus 86.2%). The F1 scores were very similar (89.5% for the lightweight CNN and 88.9% for ResNet50). The best-performing ResNet50 model attained 88.7% accuracy, 91.3% sensitivity, 86.0% specificity, and 86.9% precision. This improvement in sensitivity, however, came at the cost of reduced specificity and substantially greater computational demands. Consequently, the lightweight CNN offered a more balanced performance profile, particularly in reducing false positives, an essential factor in minimizing unnecessary CBCT referrals. These results support emerging evidence that task-specific, shallow networks can match the performance of deeper, more complex architectures when the diagnostic target is localized and well-defined (19).

Comparison with previous literature further clarifies the present results. Fukuda et al. (13) trained AlexNet, GoogLeNet, and VGG16 models for similar M3–IAC classification tasks and reported high accuracy; however, their approach required manual cropping and involved substantially larger model sizes, which pose challenges for clinical use. Similarly, Zhu et al. (4) utilized YOLOv4 for classification and achieved competitive results, yet YOLOv4 requires significantly greater GPU resources and longer inference times, limiting its feasibility in typical dental practice settings. The lightweight CNN developed in this study delivered comparable performance with drastically reduced computational demands, making it more suitable for chairside integration, incorporation into dental software systems, and use in resource-constrained environments.

The outcomes of this study suggest that model depth may not be the critical factor in tasks involving well-localized, high-contrast anatomical structures like

the M3 root and IAC path. Similarly, more shallow or less complex convolutional networks have achieved performance comparable to deeper architectures in some medical imaging classification tasks, indicating that simpler models can suffice when the diagnostic target is localized and well-defined. For example, comparisons across multiple CNN architectures on chest radiograph classification showed that shallower networks can achieve similar classification performance to deeper models, often with shorter training times and reduced computational burden, suggesting that extreme depth is not always necessary for high performance on medical imaging tasks with clearly delineated targets (20).

From a clinical standpoint, automated identification of cases where the M3 may contact or approximate the IAC could help standardize, and potentially improve the decision-making process regarding the need for CBCT imaging. Panoramic radiography remains the first-line modality for third molar assessment due to its low radiation dose and broad availability (4), yet its interpretation may vary among clinicians and between experience levels (9). A system capable of achieving sensitivities above 90% and AUC values approaching 0.97 may assist clinicians in triaging high-risk cases, reducing reliance on subjective evaluation of radiographic signs. This could ultimately enhance early detection of high-risk cases and potentially reduce inferior alveolar nerve injury, one of the most significant postoperative complications of third molar extraction (8).

Despite the promising findings, several limitations should be considered. The primary limitation is the lack of CBCT-verified ground truth, which restricts the assessment to radiographic proximity rather than true anatomical contact. Although expert radiologist interpretation is widely accepted (4, 15), panoramic radiographic signs may not always accurately represent three-dimensional relationships (21). Additionally, although the dataset of 765 half-panoramic ROIs was sufficient for good performance, deep learning models generally achieve better results with larger and more diverse training data. Furthermore, despite standardization efforts during image preprocessing, device-related variability, image artifacts, and exposure differences could have affected model learning and generalizability. The automated cropping system also exhibited decreased reliability in cases of extreme horizontal M3 angulation ($\geq 90^\circ$). Such teeth were excluded from this study because their roots extended beyond the bounding box region, highlighting the

difficulty of capturing extreme angulations on panoramic radiographs (22). Future research could explore adaptive cropping techniques, deformable bounding-box models, or multi-view data augmentation to improve localization and classification performance in these challenging cases. Further studies with more diverse datasets and validation results against CBCT-confirmed ground truth are suggested to improve the consistency of risk assessments for third molar surgery.

Conclusions

In this study, a fully automated workflow was developed to detect mandibular third molars (M3) and assess their relationship with the inferior alveolar canal (IAC) on panoramic radiographs. The detection model achieved perfect localization accuracy (100%) with a mean intersection-over-union (IoU) of 87.9%. The classification model (lightweight CNN) attained up to 90.4% accuracy, 90.0% sensitivity, and 91.1% specificity in classifying M3–IAC contact. Lightweight CNN performed comparably to the larger ResNet50 benchmark, while requiring substantially fewer computational resources. These findings support the use of compact, task-specific architectures for automated assessment of the M3–IAC relationship on panoramic radiographs.

Acknowledgements

This work was supported by the Dental Caries Prevention Research Center of Qazvin University of Medical Sciences.

Conflicts of interest

The authors declare no conflict of interest.

Ethical considerations

All procedures were conducted in accordance with relevant institutional and international ethical standards. The study protocol was approved by the ethics committee of Qazvin University of Medical Sciences (IR.QUMS.REC.1400.212). All panoramic radiographs were anonymized, and personal identifiers, including name, age, and gender, were removed before analysis.

Author contributions

A.A.K. and A.S.V. contributed to the study's conceptualization, design, and supervision, and revised the manuscript. M.T and A.S.V. participated in data collection and analysis. A.L. participated in data

collection and interpretation and prepared the manuscript draft. All authors have read and approved the final manuscript.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

1. Dhafar W, Mandura R, Dafar A, Alghamdi M. Reasons for third molars extraction by different health care providers. *Int J Pharm Res Allied Sci* 2020;9(2-2020):189-194.
2. Yang HM, Woo YJ, Won SY, Kim DH, Hu KS, Kim HJ. Course and distribution of the lingual nerve in the ventral tongue region: anatomical considerations for frenectomy. *J Craniofac Surg* 2009;20(5):1359-1363.
3. Huang C, Zhou C, Xu M, Zou D. Risk factors for lingual plate fracture during mandibular third molar extraction. *Clin Oral Investig* 2020;24(11):4133-4142.
4. Atieh MA. Diagnostic accuracy of panoramic radiography in determining relationship between inferior alveolar nerve and mandibular third molar. *J Oral Maxillofac Surg* 2010;68(1):74-82.
5. Bailey E, Kashbour W, Shah N, Worthington HV, Renton TF, Coulthard P. Surgical techniques for the removal of mandibular wisdom teeth. *Cochrane Database Syst Rev* 2020;7(7):Cd004345.
6. Cederhag J, Truedsson A, Alstergren P, Shi XQ, Hellén-Halme K. Radiographic imaging in relation to the mandibular third molar: a survey among oral surgeons in Sweden. *Clin Oral Investig* 2022;26(2):2073-2083.
7. Eshghpour M, Labafchi A, Samieirad S, Hosseini Abrishami M, Nodehi E, Rashid Javan A. Does the Presence of Impacted Mandibular Third Molars Increase the Risk of Bad Split Incidence During Bilateral Sagittal Split Osteotomy? *World J Plast Surg* 2021;10(1):37-42.
8. Kubota S, Imai T, Nakazawa M, Uzawa N. Risk stratification against inferior alveolar nerve injury after lower third molar extraction by scoring on cone-beam computed tomography image. *Odontology* 2020;108(1):124-132.
9. Uzun C, Sumer AP, Sumer M. Assessment of the reliability of radiographic signs on panoramic radiographs to determine the relationship between mandibular third molars and the inferior alveolar canal. *Oral Surg Oral Med Oral Pathol Oral Radiol* 2020;129(3):260-271.
10. Bardideh E, Lal Alizadeh F, Amiri M, Ghorbani M. Designing an artificial intelligence system for dental occlusion classification using intraoral photographs: A comparative analysis between artificial intelligence-based and clinical diagnoses. *Am J Orthod Dentofacial Orthop* 2024;166(2):125-137.
11. Majaj NJ, Pelli DG. Deep learning-Using machine learning to study biological vision. *J Vis* 2018;18(13):2.
12. Moran M, Faria M, Giraldi G, Bastos L, Oliveira L, Conci A. Classification of Approximal Caries in Bitewing Radiographs Using Convolutional Neural Networks. *Sensors (Basel)* 2021;21(15).
13. Fukuda M, Aiji Y, Kise Y, Nozawa M, Kuwada C, Funakoshi T, et al. Comparison of 3 deep learning neural networks for classifying the relationship between the mandibular third molar and the mandibular canal on panoramic radiographs. *Oral Surg Oral Med Oral Pathol Oral Radiol* 2020;130(3):336-343.
14. Zhu T, Chen D, Wu F, Zhu F, Zhu H. Artificial Intelligence Model to Detect Real Contact Relationship between Mandibular Third Molars and Inferior Alveolar Nerve Based on Panoramic Radiographs. *Diagnostics (Basel)* 2021;11(9).
15. Haddad Z, Khorasani M, Bakhshi M, Tofangchiha M, Shalli Z. Radiographic Position of Impacted Mandibular Third Molars and Their Association with Pathological Conditions. *Int J Dent* 2021;2021:8841297.
16. Tan M, Pang R, Le QV, editors. Efficientdet: Scalable and efficient object detection. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*; 2020.
17. Lin T-Y, Maire M, Belongie S, Hays J, Perona P, Ramanan D, et al., editors. Microsoft coco: Common objects in context. *European conference on computer vision*; 2014: Springer.
18. He K, Zhang X, Ren S, Sun J, editors. Deep residual learning for image recognition. *Proceedings of the IEEE conference on computer vision and pattern recognition*; 2016.
19. Khalkhali V, Azim SM, Dehzangi I. ExShall-CNN: An Explainable Shallow Convolutional Neural Network for Medical Image Segmentation. *Mach Learn Knowl Extr* 2025;7(1):19.
20. Bressen KK, Adams LC, Erxleben C, Hamm B, Niehues SM, Vahldiek JL. Comparing different deep learning architectures for classification of chest radiographs. *Scientific reports. Sci Rep* 2020;10(1):13590.
21. y Baena RR, Beltrami R, Tagliabo A, Rizzo S, Lupi S-M. Differences between panoramic and Cone Beam-CT in the surgical evaluation of lower third molars. *J Clin Exp Dent* 2017;9(2):e259.
22. Miloro M, DaBell J. Radiographic proximity of the mandibular third molar to the inferior alveolar canal. *Oral Surg Oral Med Oral Pathol Oral Radiol Endod.* 2005;100(5):545-549.