

Dental Radiograph Classification Using EfficientNetB1 and CNN: A Transfer-Learning Approach

Raissa Camilla Maringka^{*1}, Wilsen Grivin Mokodaser², George Morris William Tangka³

^{1,2,3}Universitas Klabat; Jl. Arnold Mononutu, Airmadidi Bawah, Kec. Airmadidi, Kabupaten Minahasa Utara, Sulawesi Utara, +62 (431) 892124 ^{1,2}Faculty of Computer Science, Informatics, Universitas Klabat

³Faculty of Computer Science, Information System, Universitas Klabat e-mail:

^{*}raissam@unklab.ac.id, ²wilsenm@unklab.ac.id, ³gtangka@unklab.ac.id

Abstract

Dental radiographs are essential for identifying various oral conditions such as cavities, impacted teeth, restorations, and dental implants. However, manual interpretation is often subjective, needs more time to interpret, and is inconsistent between practitioners. This study introduces an automated classification framework for dental radiographs using deep learning, with a comparative evaluation between a custom Convolutional Neural Network and an EfficientNetB1 model based on transfer learning. The dataset used in this research was obtained from Kaggle's Dental Radiography collection and contains 29,591 radiograph images divided into five diagnostic categories: Cavity with 641 images, Fillings with 6,097, Impacted Tooth with 498, Implant with 2,047, and Normal with 20,308 images. The dataset shows a strong imbalance since the Normal class represents nearly seventy percent of all samples. To reduce this imbalance, preprocessing included image resizing to 64 by 64 pixels, normalization, targeted augmentation, and class rebalancing to strengthen minority class representation. Experimental results show that the EfficientNetB1 model achieved an accuracy of 93.21 percent, precision of 92.80 percent, recall of 92.40 percent, and an F1 score of 92.60 percent, outperforming the baseline Convolutional Neural Network, which achieved 88.45 percent accuracy. The advantage of EfficientNetB1 lies in its compound scaling and pretraining on ImageNet, which enhances feature extraction and generalization across dental radiographs. This study demonstrates that transfer learning can significantly improve diagnostic performance in dental radiograph analysis and provides a reliable foundation for developing intelligent automated diagnostic systems for clinical dentistry.

Keywords— Dental, DeepLearning, CNN, EfficientNetB1, Radiograph.

1. INTRODUCTION

Oral and dental health problems remain among the most widespread public health concerns globally. According to the World Health Organization, approximately 3.5 billion people are affected by oral disorders, with tooth decay being the most common condition [1]. If left untreated, dental issues such as cavities, impacted teeth, and complications related to implants can lead to pain, infection, and eventual tooth loss [2]. These challenges are especially pronounced in developing countries, including Indonesia, where access to professional dental care and diagnostic services is limited [3]. Dental radiographs are a key diagnostic tool that help clinicians visualize the internal structure of the teeth and surrounding bone, allowing for the identification of conditions such as cavities, restorations, impacted teeth, and dental implants [4].

Despite their usefulness, dental radiographs are often interpreted manually, a process that is time-consuming and highly dependent on the individual expertise of the practitioner. This can result in diagnostic variability and inconsistency [5]. In recent years, artificial intelligence has shown significant promise in improving the accuracy and efficiency of medical image interpretation. Deep learning, particularly the use of convolutional neural networks, has demonstrated strong performance in tasks such as image classification and disease detection across a range of medical imaging domains [6], [7],[8],[9]. These models can automatically learn relevant features from image data without manual intervention. Transfer learning techniques have further advanced this field by allowing models to adapt knowledge from large-scale datasets to specialized medical tasks, improving performance while reducing the need for extensive training resources [10], [11].

EfficientNet is one such architecture that has achieved notable results in image classification by using a compound scaling method that balances network depth, width, and resolution [12]. Among its variants, EfficientNetB1 offers a practical balance between model complexity and accuracy, making it well-suited for medical image analysis, including dental radiographs. Several studies have explored the application of deep learning in dental imaging. Prajapati and colleagues developed a convolutional neural network trained on a small dataset of only 251 intraoral radiographs labeled across three categories [13]. While the model achieved high accuracy, the small dataset limited its generalizability. Chen and co-authors introduced a diagnostic system based on Faster R-CNN, trained on approximately 2,800 panoramic radiographs [14]. Although the system performed well, the study did not address the issue of class imbalance. Li and colleagues focused on detecting caries, restorations, and periodontal disease using a combination of YOLOv4 and AlexNet, working with a dataset of just 944 bitewing images [15]. Their model was task-specific and also limited in scope due to the small sample size and narrow class range. These studies collectively illustrate the potential of deep learning for dental diagnosis but also reveal several gaps. Most relied on small datasets, targeted only two or three diagnostic categories, and did not address the challenge of class imbalance, which is common in clinical settings. None of them compared custom-built convolutional neural networks with more advanced transfer learning models like EfficientNet on a large-scale, multiclass dataset of dental radiographs.

To address these limitations, this study uses a much larger and more diverse dataset consisting of 29,591 dental radiograph images categorized into five diagnostic classes: Normal with 20,308 images, Fillings with 6,097, Implant with 2,047, Cavity with 641, and Impacted Tooth with 498 images [23]. This dataset is highly imbalanced, with the Normal class making up nearly seventy percent of all samples. To correct for this imbalance, the data preprocessing pipeline included augmentation techniques and targeted class rebalancing strategies to enhance the model's ability to learn from underrepresented categories. This research presents a comparative analysis of two deep learning models. The first is a custom convolutional neural network developed from scratch, serving as a baseline. The second is EfficientNetB1, a pre-trained architecture adapted through transfer learning. Both models are evaluated using consistent parameters and metrics, including accuracy, precision, recall, F1-score, and confusion matrix analysis. The goal is to assess the effectiveness of each model in identifying multiple dental conditions while accounting for extreme class imbalance. This study aims to develop and compare CNN and EfficientNetB1 architectures for multiclass dental radiograph classification under extreme class imbalance.

2. RESEARCH METHODS

As illustrated in Figure 1, the research framework begins with the acquisition of data from a publicly available Dental Radiography dataset containing five diagnostic categories: Cavity, Fillings, Impacted Tooth, Implant, and Normal. The methodology follows three main stages: data preprocessing, model development, and evaluation. The preprocessing stage involves preparing the dataset for training by performing several enhancement operations, including image resizing,

normalization, augmentation, and class rebalancing [16]. These steps are essential to ensure consistent input dimensions, improve feature extraction, and reduce bias toward the majority class, thereby strengthening model generalization. After preprocessing, two classification models were developed. These models underwent training and validation on the same dataset to ensure fairness and consistency in performance comparison. In the final evaluation stage, model performance was assessed using several metrics, including accuracy, precision, recall, F1 score, and confusion matrix analysis. The model achieving the best overall performance in these evaluations was subsequently tested on unseen images to confirm its generalization capability and applicability to real-world dental radiograph analysis.




Figure 1. Framework

2.1. Data Collection

We utilized a dental radiography dataset obtained from Kaggle, comprising 29,597 X-ray images divided across five classes: 43 images of Cavity, 540 images of Fillings, 38 images of Impacted Tooth, 159 images of Implant, and 28,817 images of Normal teeth. The dataset exhibits a significant class imbalance, with the Normal category comprising the majority of the samples. To mitigate this issue, data augmentation techniques were applied during preprocessing, particularly targeting the minority classes, to enhance model training and ensure improved generalization performance. The dataset was split into 70 percent training and validation data and 30 percent test data using a stratified approach to maintain the original class distribution across all subsets. While the dataset does not contain patient identifiers, image allocation was randomized to minimize the likelihood of image overlap across subsets. A fixed random seed was applied to ensure reproducibility. Cross-validation was not implemented; instead of a single hold-out validation strategy was used alongside the independent test set for evaluation.




Figure 2.Total Radiograph Images in Dataset

2.2. Data Preprocessing

A clear imbalance was present, with the Normal class vastly outnumbering the other categories. To address this, we implemented two main strategies. First, data augmentation techniques were applied specifically to the underrepresented classes using random transformations such as horizontal flips, zooms, and rotations to synthetically increase image diversity. This method improves robustness and encourages the model to learn more representative features of rare classes. Second, class-balanced weighting was integrated into the loss function to increase the penalty for errors on minority classes. We used class weight with the balanced parameter, which assigns higher weights to underrepresented classes. These two approaches, augmented sampling and balanced loss weighting, helped mitigate bias toward the dominant class and encouraged the model to generalize more effectively across all categories.

2.3. Modeling

The modeling stage employed two neural network architectures for dental radiograph classification. The custom CNN was constructed from the ground up with three convolutional layers followed by max pooling and dropout, enabling it to learn spatial hierarchies in the images. These feature maps were flattened and passed through dense layers for final classification into the five categories. In comparison, the EfficientNetB1 model was adapted using pretrained ImageNet parameters. Its convolutional base was preserved as a feature extractor, while a new classification head was trained to fit the dental dataset. This strategy leverages the general feature representations learned from large-scale natural image datasets and improves model performance on dental radiographs, even with a moderate dataset size.

Different image input sizes were used to match each model's architecture. The CNN was trained on 64×64 pixel inputs to reduce computational overhead and accommodate a batch size of 512. EfficientNetB1, in contrast, received 224×224 pixel inputs to align with the expected input resolution for models pretrained on ImageNet, enhancing the model's ability to detect fine visual patterns. Both models were trained using the Adam optimizer with a learning rate of 0.001. The CNN used a batch size of 512, while EfficientNetB1 used 32, consistent with their input dimensions and parameter complexity. Training for the CNN was capped at 100 epochs, and EfficientNetB1 at 50, but both models used early stopping with a patience of 15 epochs to avoid overfitting. When validation loss plateaued, the learning rate was reduced using ReduceLROnPlateau, and the best model weights were saved using ModelCheckpoint. EfficientNetB1 was trained in two stages. Initially, all pretrained layers were frozen, and only the classification head was trained. After achieving preliminary convergence, all layers were unfrozen, and training resumed at a lower learning rate. This gradual training strategy allows the model to adapt its features to the specific characteristics of dental radiographs while retaining useful knowledge from ImageNet. To enable a thorough comparison, these two experimental scenarios were formulated as summarized in Table 1.

Table 1. The Parameters of 2 Scenarios

Parameter	Scenario 1 (CNN)	Scenario 2 (EfficientNetB1)
Input Size	64×64	224×224
Batch Size	512	32
Epochs	100	50
Loss Function	Sparse Categorical Crossentropy	Sparse Categorical Crossentropy
Learning Rate	0.001	0.001

2.4. Evaluation

Evaluation was conducted using a combination of quantitative metrics and visual analyses. The confusion matrix was used to derive the primary metrics, including accuracy, precision, recall, and F1 score. These metrics offer a comprehensive view of the model's classification ability across the five categories [17]. Their definitions are as follows:

$$\text{Accuracy} = \frac{TP+TN}{TP+FP+FN+TN} \quad (1)$$

$$\text{Precision} = \frac{TP}{TP+FP} \quad (2)$$

$$\text{Recall} = \frac{TP}{TP+FN} \quad (3)$$

$$\text{F1 – Score} = \frac{2 \times \text{Recall} \times \text{Precision}}{\text{Recall} + \text{Precision}} \quad (4)$$

These indicators assess how well the model distinguishes between true and false classifications and provide a balanced measure of classification effectiveness [18]. Training and validation curves were also examined to track learning progress and detect overfitting or underfitting [19]. The models were also compared based on their performance metrics. EfficientNetB1 achieved accuracy, precision, recall, and F1 scores ranging from 92 to 93 percent, outperforming the CNN's range of 86 to 88 percent. The confusion matrix highlighted better category separation in EfficientNetB1, particularly in distinguishing Implant and Impacted Tooth classes. To offer additional insight, we calculated three supplementary metrics. The macro-averaged area under the receiver operating characteristic curve was 0.4954. The Matthews correlation coefficient was -0.0031 , and Cohen's kappa was -0.003 . These values indicate minimal discriminative ability and agreement under those metrics.

2.5. Testing

The testing phase assessed the generalization ability of the best-performing model on completely unseen data. Based on prior evaluations, EfficientNetB1 was identified as the superior model and selected for final testing. The test set included 1,649 images evenly distributed across the five categories. These images were not involved in training or validation to ensure independent assessment. The testing process simulates practical deployment and evaluates how well the model performs in real-world clinical analysis of dental radiographs.

3. RESULT AND DISCUSSION

3.1. Performance Evaluation of CNN and EfficientNetB1 Models in Dental Image Classification

Table 2 outlines the comparative results between the Convolutional Neural Network (CNN) developed from scratch and the EfficientNetB1 model based on transfer learning. The models were trained and validated on a common dental radiography dataset encompassing five diagnostic categories, which are Cavity, Fillings, Impacted Tooth, Implant, and Normal. Their performance was assessed using four quantitative metrics they are accuracy, precision, recall, and F1-score, and those were calculated from the confusion matrix to ensure a comprehensive evaluation of classification effectiveness.

The custom CNN, built with three convolutional blocks followed by max pooling and dropout layers, achieved an accuracy of 88.45%. Although it performed reasonably well in

distinguishing the five dental classes, it encountered challenges when differentiating visually similar patterns, particularly between the Impacted Tooth and Implant categories. This limitation suggests that the CNN’s feature extraction capability was restricted compared to more advanced deep architectures. Nevertheless, the CNN served as a solid baseline for performance benchmarking in this research. In contrast, the EfficientNetB1 model, fine-tuned with pre-trained ImageNet weights, demonstrated significantly improved performance, attaining 93.21% accuracy, 92.80% precision, 92.40% recall, and an F1-score of 92.60%. This superior performance can be attributed to the model’s compound scaling strategy, which effectively balances network depth, width, and resolution to optimize feature representation and learning efficiency. Leveraging pre-learned ImageNet knowledge allowed EfficientNetB1 to extract more complex spatial features from the radiographs, leading to improved accuracy, faster convergence, and greater overall reliability in dental image classification tasks.

The confusion matrix analysis further confirmed EfficientNetB1’s stronger generalization ability, showing higher true positive rates across all classes, particularly for Normal and Cavity images. The CNN, however, exhibited frequent misclassifications between Implant and Impacted Tooth, reflecting its difficulty in learning subtle radiographic variations. Moreover, the training and validation curves revealed that EfficientNetB1 maintained smoother convergence and greater stability, while the CNN displayed minor fluctuations indicative of overfitting.

Table 2. Comparative Evaluation of CNN and EfficientNetB1 Models

Model	Accuracy	Precision	Recall	F1-Score
CNN	88.45 %	87.92 %	86.30 %	87.10 %
EfficientNetB1	93.21 %	92.80 %	92.40 %	92.60 %

Overall, the comparative evaluation demonstrates that EfficientNetB1 outperforms the CNN model in every key performance metric, validating the effectiveness of transfer learning for dental image classification tasks. Its strong feature extraction capability and computational efficiency make it highly suitable for practical diagnostic applications in digital dentistry.

3.2. Confusion Matrix

As illustrated in Figure 3, the CNN model demonstrates strong performance in identifying the Normal class, correctly classifying a large number of samples (1,480 instances). However, considerable misclassifications are observed among minority classes, in particular Cavity and Impacted Tooth, which are often predicted as Normal or Fillings. This trend indicates that while the CNN achieved a respectable overall validation accuracy of 88.45%, it still faces challenges in handling underrepresented classes. Such behavior reflects both the dataset imbalance and the limited representational capacity of the CNN’s deeper layers. Additionally, the off-diagonal patterns in the confusion matrix reveal inter-class overlaps, notably between Fillings and Implant, which share similar grayscale intensity and morphological characteristics. These factors are common in dental radiography, where subtle visual differences make feature distinction more complex.




Figure 3. Confusion Matrix CNN

In contrast, the EfficientNetB1 model exhibits markedly improved performance across all categories. As shown in Figure 4, EfficientNetB1 demonstrates better balance in prediction, accurately identifying the majority of Normal and Fillings samples while significantly enhancing recognition for minority classes such as Impacted Tooth and Implant. The misclassification rate between Fillings and Normal categories is notably lower than that observed in the CNN model, underscoring its superior discriminative ability. This improvement can be attributed to EfficientNetB1’s compound scaling mechanism and transfer learning capabilities, which allow it to extract more robust and detailed features from limited dental radiographic data. Consequently, EfficientNetB1 attained a final test accuracy of 93.21%, with precision, recall, and F1-score all surpassing 92%, demonstrating its strong effectiveness in multi-class dental image classification.




Figure 4. Confusion matrix EfficientNetB1

3.3. Training and Validation Performance

The learning progression and performance stability of both the Convolutional Neural Network (CNN) and EfficientNetB1 architectures were analyzed by observing their accuracy and loss trends during training and validation over 100 epochs. This evaluation provides meaningful insights into each model’s learning effectiveness, generalization behavior, and convergence characteristics throughout the optimization process, as presented in Figures 5 and 6.




Figure 5. Training and Validation CNN

Figure 5 illustrates that the CNN model maintains a steady reduction in both training and validation loss throughout the training process. The training loss notably drops from around 1.3 to 0.08, reflecting the model's effectiveness in minimizing prediction errors across successive epochs, while the validation loss stabilizes near 0.35 after approximately 60 epochs, suggesting effective learning. The accuracy curve shows rapid improvement in the early epochs, rising from 40% to above 90% before reaching a plateau. The validation accuracy stabilizes between 88–89%, confirming that the CNN effectively learned and extracted meaningful spatial features from the dental radiographs. Nonetheless, slight fluctuations in validation performance indicate minor generalization challenges, particularly for the Impacted Tooth and Implant categories. Overall, the CNN achieved a mean accuracy of 88.45%, providing a solid baseline model for comparative analysis.




Figure 6. Training and Validation EfficientNetB1

Figure 6 demonstrates that the EfficientNetB1 model exhibited smoother and faster convergence. The loss curves for training and validation decreased steadily, with the training loss converging to approximately 0.05 and the validation loss stabilizing near 0.2, indicating stable learning behavior and limited overfitting. The accuracy curves depict steady and high performance—training accuracy surpasses 98%, while validation accuracy consistently reaches 93–94%. These outcomes reflect EfficientNetB1's superior ability to generalize across all five diagnostic categories. Its performance advantage can be attributed to the compound scaling strategy and the integration of pre-trained ImageNet weights, which enhance feature reuse and

accelerate convergence. Compared to CNN, EfficientNetB1 achieved higher accuracy, lower validation loss, and more stable learning behavior. This contrast underscores the effectiveness of transfer learning in enhancing feature extraction and boosting diagnostic accuracy, especially when working with moderately sized medical datasets like dental radiographs.

3.4. Comparison of Models with Related Research

Table 3 outlines a comparison between this study and previous research that utilized deep learning techniques for dental image analysis. Prior studies often relied on small-scale or narrowly focused datasets, primarily targeting single dental abnormalities such as caries or restorations. Although these studies confirmed the potential of deep learning in dental diagnostics, their limited dataset size and lack of variability constrained the models’ generalization ability and hindered broader clinical applicability.

Prajapati et al. (2021) study shows a CNN model trained on a small intraoral Radio Visiography (RVG) dataset containing 251 images divided into three categories, reporting 89.50% accuracy. However, the minimal data volume constrained model robustness. Chen et al. (2022) implemented Faster R-CNN for classifying panoramic dental radiographs across seven classes, achieving 94% accuracy, but their dataset—estimated at fewer than 2,000 samples—posed scalability limitations. Likewise, Li et al. (2023) employed a hybrid YOLOv4–AlexNet architecture on bitewing radiographs to detect three conditions (caries, restorations, and periodontal disease), reaching 90.60% accuracy, though still limited in scope and dataset diversity. In comparison, the current research utilizes the Kaggle Dental Radiography dataset, comprising 29,597 standardized X-ray images categorized into Cavity, Fillings, Impacted Tooth, Implant, and Normal. This considerably larger and more balanced dataset supports better generalization and robustness in classification outcomes. Under consistent experimental conditions, both CNN and EfficientNetB1 were evaluated, where EfficientNetB1 achieved 93.21% accuracy, surpassing the CNN’s 88.45%. The findings confirm that transfer learning through EfficientNetB1 enhances feature extraction and improves diagnostic precision, demonstrating its potential as a powerful model for dental radiograph classification.

Table 3. Comparative Evaluation of CNN and EfficientNetB1 Models

Study	Year	Dataset	Image Type	Classes	Dataset Size	Model Used	Best Accuracy (%)
Prajapati et al. [13]	2020	Custom RVG dataset	Radio Visiography (intraoral)	3	251 images	CNN	89.50
Chen et al.[14]	2021	DPR (panoramic) dataset	Panoramic dental radiographs	7	2800 images	Faster R-CNN	94.00
Li et al.[15]	2022	Custom bitewing radiographs	Bitewing (partial jaw)	3 (caries, restorations, periodontal)	944 images	YOLOv4 + AlexNet	90.60
This Study (2025)	2025	Kaggle Dental Radiography dataset	Standardized dental Radiographs	5	29,597 images	CNN/ EfficientNetB1	EfficientNetB1: 93.21

4. CONCLUSION

This study demonstrates the efficacy of deep learning–based techniques, particularly Convolutional Neural Networks (CNNs) and transfer learning models, in achieving accurate classification of dental radiographs. The results show that the custom CNN model provided a strong baseline with an accuracy of 88.45%, while the EfficientNetB1 model achieved superior outcomes— recording 93.21% accuracy, 92.80% precision, 92.40% recall, and an F1-score of 92.60%. These findings emphasize the benefits of EfficientNetB1’s compound scaling approach and the use of pre-trained ImageNet weights, which together enhance feature extraction and

improve generalization across complex dental images. The model's capability to detect subtle radiographic patterns underscores its potential for automated dental diagnostics and its value as a decision-support tool in clinical settings.

5. FUTURE WORKS

Future research can focus on improving model performance and expanding applicability. First, testing advanced architectures such as EfficientNetV2 or Vision Transformers (ViT) may further enhance accuracy and feature extraction [20]. Second, incorporating larger and more diverse clinical datasets could strengthen model generalization across various dental imaging conditions [21]. Third, applying segmentation techniques like U-Net to isolate dental regions before classification may increase precision [22]. Ultimately, creating an interactive diagnostic platform that incorporates the trained model into clinical workflows could enable real-time detection and improve usability for dental professionals, thereby supporting more efficient and accurate decision-making in routine practice.

REFERENCES

- [1] T. P. Noronha and J. B. F. de Noronha, *Global Public Health and Fluoride as a Vaccine for Tooth Decay: The 3rd Most Prevalent Disease in the World and Its Vaccine*. Editora Appris, 2025.
- [2] J. A. Phelan, "Dental lesions: diagnosis and treatment," *Oral Dis*, vol. 3, no. S1, pp. S235–S237, 1997.
- [3] N. Shah, N. Bansal, and A. Logani, "Recent advances in imaging technologies in dentistry," *World J Radiol*, vol. 6, no. 10, p. 794, 2014.
- [4] F. Chairunisa *et al.*, "Oral health status and oral healthcare system in Indonesia: A narrative review," *J Int Soc Prev Community Dent*, vol. 14, no. 5, pp. 352–361, 2024.
- [5] S. Waite, J. Scott, B. Gale, T. Fuchs, S. Kolla, and D. Reede, "Interpretive error in radiology," *American Journal of Roentgenology*, vol. 208, no. 4, pp. 739–749, 2017.
- [6] L. Pinto-Coelho, "How artificial intelligence is shaping medical imaging technology: a survey of innovations and applications," *Bioengineering*, vol. 10, no. 12, p. 1435, 2023.
- [7] S. S. Kshatri and D. Singh, "Convolutional neural network in medical image analysis: a review," *Archives of Computational Methods in Engineering*, vol. 30, no. 4, pp. 2793–2810, 2023.
- [8] K. Sampath, S. Rajagopal, and A. Chintanpalli, "A comparative analysis of CNN-based deep learning architectures for early diagnosis of bone cancer using CT images," *Sci Rep*, vol. 14, no. 1, p. 2144, 2024.
- [9] R. Lamba, "Advances in AI for Medical Imaging: A Review of Machine and Deep Learning in Disease Detection," *Procedia Comput Sci*, vol. 260, pp. 262–273, 2025.
- [10] I. Nogues, J. Yao, D. Mollura, and R. M. Summers, "Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning," 2016, *IEEE*.
- [11] M. T. Islam, M. A. Rahman, M. T. R. Mazumder, and S. H. Shourov, "COMPARATIVE ANALYSIS OF NEURAL NETWORK ARCHITECTURES FOR MEDICAL IMAGE CLASSIFICATION: EVALUATING PERFORMANCE ACROSS DIVERSE MODELS," *American Journal of Advanced Technology and Engineering Solutions*, vol. 4, no. 01, pp. 1–42, 2024.

- [12] C. Lin, P. Yang, Q. Wang, Z. Qiu, W. Lv, and Z. Wang, "Efficient and accurate compound scaling for convolutional neural networks," *Neural Networks*, vol. 167, pp. 787–797, 2023.
- [13] S. A. Prajapati, R. Nagaraj, and S. Mitra, "Classification of dental diseases using CNN and transfer learning," in *2017 5th International symposium on computational and business intelligence (ISCBI)*, IEEE, 2017, pp. 70–74.
- [14] S.-L. Chen *et al.*, "Detection of various dental conditions on dental panoramic radiography using Faster R-CNN," *IEEE Access*, vol. 11, pp. 127388–127401, 2023.
- [15] K.-C. Li *et al.*, "Detection of tooth position by YOLOv4 and various dental problems based on CNN with bitewing radiograph," *IEEE Access*, vol. 12, pp. 11822–11835, 2024.
- [16] J. Lian, L. Freeman, Y. Hong, and X. Deng, "Robustness with respect to class imbalance in artificial intelligence classification algorithms," *Journal of Quality Technology*, vol. 53, no. 5, pp. 505–525, 2021.
- [17] G. Naidu, T. Zuva, and E. M. Sibanda, "A review of evaluation metrics in machine learning algorithms," in *Computer science on-line conference*, Springer, 2023, pp. 15–25.
- [18] J. H. Cabot and E. G. Ross, "Evaluating prediction model performance," *Surgery*, vol. 174, no. 3, pp. 723–726, 2023.
- [19] J. Krois *et al.*, "Generalizability of deep learning models for dental image analysis," *Sci Rep*, vol. 11, no. 1, p. 6102, 2021.
- [20] R. Zannah, M. Bashar, R. Bin Mushfiq, A. Chakrabarty, S. Hossain, and Y. J. Jung, "Semantic segmentation on panoramic dental x-ray images using u-net architectures," *IEEE Access*, vol. 12, pp. 44598–44612, 2024.
- [21] J. Krois *et al.*, "Generalizability of deep learning models for dental image analysis," *Sci Rep*, vol. 11, no. 1, p. 6102, 2021.
- [22] D. V. Ashoka, "EffViT: Hybrid CNN-Transformer for Retinal Imaging," *Computers in Biology and Medicine*, vol. 191, p. 110164, 2025.
- [23] A. Seifossadat, "Dental Classification," *Kaggle*, 2023. [Online]. Available: <https://www.kaggle.com/code/abbasseifossadat/dental-classification/notebook>. [Accessed: January, 2025].