Volume 4, Issue 3, 2026

Frontier in Medical & Health Research

ISSN: (e) 3007-1607 (p) 3007-1593

DEEP LEARNING-BASED SEGMENTATION OF DENTAL INFECTIONS IN RADIOGRAPHIC IMAGES USING U-NET

Mareena Karim[*1], Seena Irshad[2], Ali Nasir[3], Abdullah Ismail[4], Sahib Khan[5] Imran Maqsood[6]

[*1,3,4,5]*Department of Telecommunications Engineering, University of Engineering & Technology, Mardan, Pakistan*
[2]*Department of Paediatric Dentistry, Bacha Khan Dental College and Hospital, Mardan, Pakistan*
[6]*Department of Computer Software Engineering, University of Engineering & Technology, Mardan, Pakistan*

[*1]mareenakarim@uetmardan.edu.pk

DOI: https://doi.org/10.5281/zenodo.21882651

Keywords
Artificial intelligence; convolutional neural network; U-Net; image segmentation; dental radiography; dental infection; diagnostic support.

Article History
Received: 25 January 2026
Accepted: 10 March 2026
Published: 24 March 2026

Copyright @Author

Corresponding Author: *
Mareena Karim

Abstract

Clinical importance: Well-defined borders are essential in the diagnosis and treatment planning of dentition and are important for follow-up evaluation in dental X-rays. Manual segmentation is time consuming, and may be variable between observers. Objective: To build a deep learning pipeline for binary, pixel-wise dental infection region segmentation in X-rays with U-Net. The project was based on an existing annotated collection of grayscale dental X-ray images, image level annotations and binary segmentation. The data was divided in to 70% training data, 15% validation data and 15% test data. The U-Net model with 4 encoder blocks and 4 decoder blocks was implemented, along with skip connections, ReLU activation, batch normalization and dropout. The training was done on Google Colaboratory with a learning rate of 0.0001, a batch size of 16, a mixture of objectives (Dice loss and binary cross-entropy), and a planned number of 50–100 epochs with early stopping on validation loss. Results: The project documentation was archived and includes details about the model implementation, training, testing, and generation of segmentation masks separating infected and normal tissue. The source thesis, however, wasn't able to provide the number of images in the dataset, the resolution of each image, exact training time, number of outputs per image, number of images per confusion-matrix cell, and numerical dice, IoU, precision, recall, and F1. Therefore, this article is a development and proof-of-concept evaluation and is not intended to provide a diagnostic-accuracy estimate. Conclusion: The pipeline implementation shows the potential of using a compact U-Net to automate the segmentation of dental infections and suggests that comprehensive documentation of datasets and numerical result of the test sets along with external clinical validation is required prior to any diagnostic claim.

1. Introduction

A dental radiograph can be helpful in revealing pathological changes which are not evident at the time of clinical exam. Periapical, panoramic, bitewing, and 3-D imaging can help to assess the extent of caries, periodontal bone loss, periapical radiolucencies, fractures, and various other structural irregularities. One of the large computational challenges in this workflow is to assign pixels in an image to the region in which they reside, with a clinical significance so that an abnormal area can be localized, measured and compared over a period of time. Manual delineation by a dental professional is still a

Frontier in Medical & Health Research

ISSN: (e) 3007-1607 (p) 3007-1593

valuable reference technique, but it is slow, difficult to scale and varying in terms of experience and interpretation. Low contrast, noise, overlapping anatomy, restorative materials, and large variability of morphology of teeth are also technical difficulties for dental radiographs. Often, parameter tuning has to be done for each image and traditional methods such as thresholding, edge detection, region growing and contour based methods may not be applicable if the lesion boundaries are weak and irregular.

Alternatively, deep learning, which learns features directly from annotated examples. Convolutional neural networks (CNNs) have been successfully applied to dental image analysis and have the ability to learn hierarchical representations of texture, shape, and spatial context [2, 3, 17]. U-Net is especially applicable to segmentation tasks, as the contextual features are captured at the encoding side, and then localized at a high resolution through skip connections [19]. These properties make them suitable to dental lesions, because the area of interest may be a small, poorly defined portion of an image. In the final year project which this article is based on, a U-Net pipeline was used to detect the infected areas in images of dental X-rays. The model structure and data organization were described, and the data-split ratio and loss function were given, as were the optimizer, learning rate, and batch size; but the numerical outputs of the tests for diagnostic accuracy were not saved. This paper thus represents a synthesis of the evidence available, and is a transparent implementation study in that it makes a clear distinction between evidence that was documented and evidence that could not be verified.

This study aimed to develop and test a Deep Learning Segmentation Pipeline for Binary Segmentation of Dental Infection Regions in Dental Radiographic Images. The ultimate aim was to design an automated, pixel-wise segmentation method that can be used to separate infected and healthy dental structures, taking into account issues of low contrast, anatomical superimposition, variability in images captured, and the limitations of traditional image-processing algorithms. The specific goals were to (i) create a segmentation model using the U-Net network for dental radiographic images, (ii) design a pre-processing and training pipeline for annotated dental image sets, (iii) compare the performance of the segmentation framework with standard performance measures including the Dice coefficient, Intersection over Union (IoU), precision, recall, F1-score and performing confusion-matrix analysis, and (iv) investigate the potential of deep learning as a supporting technology for dental diagnosis and treatment planning. The study also hopes to be a first step towards future studies in this domain with bigger datasets, more sophisticated architectures and clinically validated models.

2. Related Work

2.1 Conventional Dental Image Segmentation

Conventional dental image segmentation methods typically segment the image using a fixed intensity threshold. Earlier approaches were based on intensity thresholds, edge operators, watershed transforms, active contours, and graph-based methods. These methods can perform well on high-contrast, well-positioned images but tend to perform poorly in the presence of overlapping teeth, faint lesions, and variability in image acquisition. They also rely on manual selection of parameters, which restricts transferability across datasets and clinical settings [1], [17].

2.2 CNNs and the U-Net Family

U-Net introduced a symmetric contracting and expanding architecture for biomedical image segmentation [19]. The contracting path extracts increasingly abstract features, while the expanding path restores spatial information. Higher-resolution features from the encoder are passed to the decoder through skip connections, which helps minimize the loss of boundary information. Later variants incorporate attention mechanisms, residual connections, recurrent refinement, denoising encoders, and multi-scale fusion to improve segmentation under noise and anatomical variability [5], [8], [11].

2.3 Applications in Dental Radiography

Deep learning methods have been applied to tooth segmentation, periodontal assessment, apical-lesion localization, restoration detection, and

Frontier in Medical & Health Research

ISSN: (e) 3007-1607 (p) 3007-1593

volumetric analysis of CBCT data. Widyaningrum et al. compared multi-label U-Net and Mask R-CNN for periodontitis-related panoramic segmentation [4]. Boussaid et al. [5] integrated attention mechanisms into Teeth U-Net to improve feature selection in low-contrast panoramic images, and Li et al. [6] achieved improved tooth segmentation using an attention U-Net. Rahman et al. used a denoised residual encoder to address radiographic noise [7], and Kumar et al. proposed a fusion of multi-scale and recurrent residual learning [8].

Hybrid approaches have also been explored. Shan et al. combined U-Net output with morphological post-processing to separate adjacent tooth instances [9]. Chang et al. combined a boundary-aware U-Net with an object-detection pipeline for instance segmentation and classification of teeth [10]. Other studies have investigated attention-gated fusion [11], apical-lesion segmentation [12], restoration segmentation [13], lightweight models [14], nnU-Net for CBCT [15], multimodal dental datasets [16], and optimization-assisted 3-D segmentation [18].

2.4 Research Gap

Although deep neural networks (DNN) are seen to perform better than hand-coded rules for image processing in the literature, architecture is not enough for clinical translation. Documented data provenance, patient-level splitting, annotation procedures, hardware and software configuration, training history, held out test metrics, uncertainty estimates, and external validation are also crucial for reliable studies. The source project successfully documented several other elements but not these reproducibility elements. The aim of this paper is therefore to make a two-fold contribution: to report the development process for U-Net and to explicitly state the evidence that is still needed before diagnostic performance can be claimed.

3. Materials and Methods

3.1 Study Design and Source Evidence

This study is a retrospective algorithm-development and proof-of-concept research project, called "Image Segmentation Based on Deep Learning Model for Identification of Dental Images" as the final year research project. Any prospective patient interventions or recruitment are not reported in the study records.

To maintain the integrity of the research, only values as reported in the source thesis were re-stated in this article, no synthetic dataset counts, diagnostic metrics, confusion-matrix values, demographic distributions, or training curves were created. Table 1 categorises the information that was retained in the thesis and the information that was not retained.

Table 1. Dataset and study information documented in the source thesis.

Parameter	Available information	Status
Image type	Grayscale dental X-ray images	Reported
Segmentation target	Identification of infected and non-infected regions	Reported
Dataset organization	images/, annotations/, and labels/ folders containing images, annotations, and pixel-wise masks	Reported
Ground-truth masks	Pixel-wise segmentation masks manually prepared for infected regions	Reported
Dataset split	Training: 70%; Validation: 15%; Testing: 15%	Reported
Total number of images	Not reported	Not reported

Frontier in Medical & Health Research

ISSN: (e) 3007-1607 (p) 3007-1593

Parameter	Available information	Status
Image resolution	Not reported	Not reported
Imaging modality details	Dental radiographic images; exact acquisition parameters not reported	Partially available
Patient demographic information	Not reported	Not reported
Dataset source/licensing	Not reported	Not reported
Annotation procedure	Annotations and masks were available; annotator details were not reported	Partially available

3.2 Methodology Workflow

The steps of the workflow were: Organisation of the datasets, feature learning through the U-Net network, Pixel-wise segmentation, Classifying infected and non-infected pixels, Evaluation of performances, and Generation of Output-mask. The segmentation model was the heart of the processing phase that stood between the input dataset and the evaluation/output components as shown in figure 1.

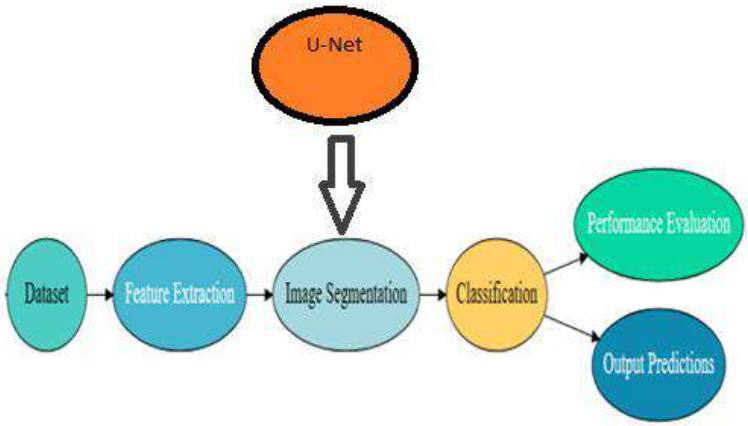


Figure 1. Project workflow for U-Net-based dental image segmentation, adapted from the study documentation

3.3 Dataset Organization and Partitioning

The dataset is split into three directories containing grayscale images, image level annotations and pixel level annotations respectively. The infected pixels were considered as the 'positive' class and healthy/background pixels were considered as the 'negative' class in segmentation masks. The entire data set was divided into training, validation and test sets in the ratio 70:15:15 as shown in Table 1.

The thesis did not specify if this was done at the image level or at the patient level. Patient-level partitioning is important, as in the future multiple images may be generated from the same patient, which may lead to multiple images being split into both the training and test sets, and hence the actual performance may be over-estimated if splitting is performed at the image-level.

3.4 U-Net Architecture

The standard U-Net concept illustrated in Figure 2 was used for the model. The encoder gradually decreased the spatial resolution and increased the level of abstraction of the features, and the

Frontier in Medical & Health Research

ISSN: (e) 3007-1607 (p) 3007-1593

decoder gradually increased the spatial resolution. Skip connections connect feature maps of the same stage of the encoder to the decoder to facilitate fine spatial detail recovery.

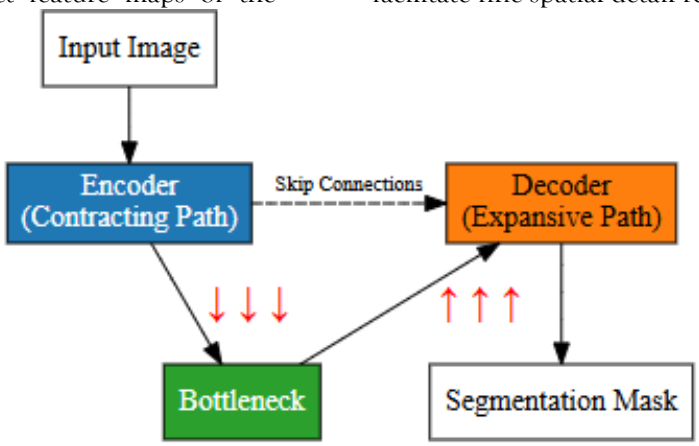


Figure 2. Conceptual U-Net architecture used in the project, showing the encoder, bottleneck, decoder, skip connections, and output segmentation mask

Table 2. Main components of the implemented U-Net.

Component	Project implementation	Purpose
Encoder	Four down-sampling blocks with convolution and max-pooling	Extract hierarchical spatial features
Bottleneck	Lowest-resolution feature representation	Integrate high-level context
Decoder	Four up-sampling blocks	Restore spatial resolution
Skip connections	Encoder features concatenated with corresponding decoder features	Preserve localization and boundary detail
Activation	ReLU	Introduce non-linearity
Normalization	Batch normalization	Stabilize and accelerate optimization
Regularization	Dropout	Reduce overfitting
Output	Binary pixel-wise segmentation mask	Separate infected from non-infected pixels

3.5 Training Configuration

The model was trained on Google Colaboratory. The thesis specified a combined Dice-loss and binary cross-entropy objective, the Adam optimizer with a learning rate of 0.0001, a batch size of 16, and early stopping based on validation loss. These settings are summarized in Table 3.

Table 3. Training configuration recorded in the study documentation.

Hyperparameter	Recorded implementation details	Status
Computing environment	Google Colaboratory	Reported

Frontier in Medical & Health Research

ISSN: (e) 3007-1607 (p) 3007-1593

Hyperparameter	Recorded implementation details	Status
Deep learning framework/version	Not reported	Not available
Hardware/GPU configuration	Not reported	Not available
Network architecture	U-Net with four down-sampling and four up-sampling blocks	Reported
Loss function	Dice Loss + Binary Cross-Entropy	Reported
Optimizer	Adam optimizer	Reported
Learning rate	0.0001	Reported
Batch size	16	Reported
Training epochs	50–100 planned epochs	Reported
Early stopping	Based on validation loss	Reported
Actual stopping epoch	Not reported	Not available
Training duration	Not reported	Not available
Random seed/reproducibility settings	Not reported	Not available

3.6 Evaluation Metrics

The proposed assessment combined overlap-based measures with pixel-level classification measures. The metrics identified in the study documentation are described in Table 4. True positives (TP) are infected pixels correctly predicted, false positives (FP) are healthy or background pixels incorrectly predicted as infected, and false negatives (FN) are infected pixels that the model failed to predict.

Table 4. Evaluation metrics defined for the held-out test set.

Metric	Definition	Interpretation
Dice coefficient	2TP / (2TP + FP + FN)	Spatial overlap between prediction and ground truth
Intersection over Union	TP / (TP + FP + FN)	Proportion of the union correctly overlapped
Precision	TP / (TP + FP)	Reliability of positive pixel predictions
Recall (sensitivity)	TP / (TP + FN)	Proportion of infected pixels detected
F1-score	2 × precision × recall / (precision + recall)	Balance between precision and recall
Confusion matrix	Counts of TP, TN, FP, and FN	Detailed pixel-level error profile

Frontier in Medical & Health Research

ISSN: (e) 3007-1607 (p) 3007-1593

3.7 Preprocessing and Post-processing

The thesis mentioned image preprocessing only in general terms and did not record the target resize dimensions, intensity-normalization method, augmentation policy, or mask-resampling procedure. Similarly, no specific morphological post-processing operation was reported. These steps are therefore not discussed in this article as fully specified experimental procedures.

3.8 Ethics and Data Governance

A pre-existing annotated dataset was used, and the archived project report did not describe prospective recruitment or intervention. However, it also did not provide the source of the original images, licensing terms, de-identification procedures, or an ethics approval/exemption reference. These elements should be verified by the authors before the dataset is submitted for journal publication.

4. Results

4.1 Implementation Outcome

According to the study documentation, the U-Net pipeline was developed on Google Colaboratory and trained on annotated dental images and tested on a held-out subset. The output masks were found to successfully distinguish infected regions from normal regions, thus lending credence to the proposed encoder-decoder approach. The project record has been stored and results are presented in table 5. The table is meant to differentiate from completed engineering work and from diagnostic-performance claims which cannot be reconstructed from the original logs and prediction files.

Table 5. Project outcomes and evidential status.

Outcome	Source-thesis finding	Interpretive status
Model construction	Four-level U-Net with skip connections, batch normalization, dropout, and ReLU was implemented	Supported
Training pipeline	Adam optimization with Dice + binary cross-entropy and early stopping was specified	Supported
Data partitioning	70:15:15 training/validation/test split was reported	Supported
Mask generation	Predicted masks were generated for infected-region segmentation	Supported qualitatively
Functional discrimination	The thesis states that infected and healthy tissue were distinguished	Qualitative project claim
Numerical Dice/IoU/F1	No values were preserved	Cannot be reported
Confusion matrix	No TP, TN, FP, or FN counts were preserved	Cannot be reported
Per-patient variability	No individual-case metric table was preserved	Cannot be reported
External validation	Not conducted or not documented	Not established

4.2 Qualitative Segmentation Findings

As per the source thesis, the trained model was able to generate output masks which could differentiate the infected from the non-infected regions, and the encoder-decoder architecture was found to be useful in addressing the issues of low contrast and overlapping structures. This is helpful to come up with a proof of concept conclusion about the feasibility of the workflow. Unfortunately, the thesis file provided for this

Frontier in Medical & Health Research

ISSN: (e) 3007-1607 (p) 3007-1593

article did not contain representative radiograph ground truth–prediction panels to allow for inspection of boundary agreement and characteristic failure modes in isolation.

4.3 Quantitative Reporting Boundary

Numerical testset measures were not found in the study documentation. Hence, in this paper the Dice coefficient, IoU, precision, recall, F1-score, pixel accuracy, confidence intervals and confusion matrix are not reported. Using other numbers than the original test predictions would give results different from the actual experiment. Hence, the study should not be reported as showing a certain diagnostic accuracy.

4.4 Reproducibility Assessment

The architectural and optimization settings presented in Tables 2 and 3 offer a good starting point for reproduction. But the information about dataset provenance/size, processing, software version, random seed, stopping (actual) epoch, hardware, training time, test prediction and annotation quality is still missing. Figure 3 shows the completeness of the documentation captured in the three domains in Tables 1, 3 and 5, with training documentation being the most completed and the outcome verification and data provenance documentation being the least completed. It would be possible to have a full quantitative evaluation if the original notebook, model checkpoint and dataset manifest were recovered.

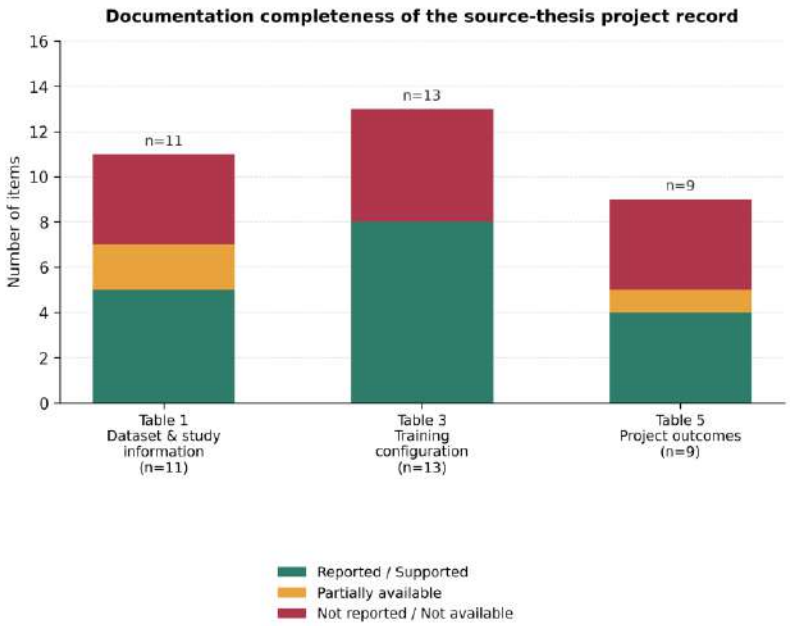


Figure 3. Documentation completeness of the source-thesis project record, summarized from the status column of Tables 1, 3, and 5.

5. Discussion

5.1 Principal Findings

The project illustrates how a standard U-Net can be developed and trained for binary segmentation of dental infection zones on an easily accessible cloud platform. Its primary practical benefit is the adoption of a well-established encoder-decoder framework, skip connections, and a class-imbalanced segmentation loss function. The work also demonstrates a problem in student and prototype research: if the metadata and test results of the experiments are not systematically recorded,

Frontier in Medical & Health Research

ISSN: (e) 3007-1607 (p) 3007-1593

the model does not result in a publishable diagnostic-accuracy study.

5.2 Relationship to Prior Work

The selected architecture is similar to previous segmentation of dental images described in Section 2. Variants of U-Net have been suggested to tackle low contrast, noise, irregular morphology and overlapping teeth [5] [8] while hybrid pipelines with morphological refinement or object detection have been suggested to enhance instance separation and classification [9] [10]. While the present architecture is simpler than these variants, this simplicity can be beneficial in a baseline study, as it makes implementation easier, and offers a clear benchmark for comparison in future studies. The reported workflow is also similar to other dental AI applications, such as periodontal assessment [4], apical-lesion segmentation [18], restoration detection [19], lightweight deployment [20] and CBCT segmentation [21]. But different from studies that show a held-out score, the evidence here is only for an implementation-level conclusion. The results cannot be compared with previously published Dice or IoU until the original model is re-tested with a documented test set.

5.3 Clinical Relevance

A robust infection-segmentation framework could help clinicians identify suspicious areas, standardize lesion measurement, and compare lesions over time. It may also serve as a pre-processing step for automated classification, treatment planning, or educational visualization. Such a framework should be used as a decision-support tool rather than a standalone diagnostic replacement, with outputs interpreted by trained dental professionals.

5.4 Strengths

There are a number of strengths in this study. Firstly, it is based on the widely used U-Net architecture, which has been shown to be effective for biomedical image segmentation. Encoder-decoder learning, skip connections, batch normalization, dropout regularization, and a combined Dice and binary cross-entropy loss are suitable solutions to handle pixel-level segmentation. A structured strategy for evaluation, which includes both an overlap-based and classification-based approach, is an additional asset. Furthermore, the study offers a clear baseline implementation which can be used as a basis for more complex future models like attention-based U-Net, residual networks and transformer-based segmentation models.

6. Conclusion

In this study, the implementation of U-Net in the last year was consolidated for binary segmentation of dental infection regions in radiographic images. The reported network was trained with Adam optimizer, learning rate of 0.0001, batch size of 16, and combined Dice and binary cross-entropy loss via a four-level encoder-decoder architecture with skip connections, batch normalization, dropout and ReLU activation function. Study documents show that training, testing, and mask generation were successful, showing the feasibility of the proposed workflow. But none of the preserved materials (datasets, pre-processing details, quantitative metrics, prediction files, or ethics and source documentation) allow for a defensible estimate of diagnostic accuracy. It should be considered therefore as a proof of concept and development study. To make the claim of clinical performance, the model would need to be rerun with all the experimental logging and independent validation.

7. Future Work

Future research is desirable to expand the proposed framework with more data obtained from various institutions and imaging systems. The current binary segmentation method could be adapted to multi-class segmentation so as to classify the dental pathologies like caries, periodontal lesions, fractures and abscesses separately. Further studies are needed on advanced architectures, such as Attention U-Net, Residual U-Net, nnU-Net, and hybrid CNN-transformer models for better feature extraction and boundary localization. Future directions also include studies on 3-D segmentation using CBCT imaging, explainable AI to help create clinician

Frontier in Medical & Health Research

ISSN: (e) 3007-1607 (p) 3007-1593

trust and prospective clinical validation to examine real-world impact.

8. Declarations

8.1 Ethics Approval and Consent to Participate

A pre-existing annotated image collection was used for this project, and no recruitment or direct patient contact was described. The original source of the dataset and the corresponding ethics approval or exemption reference were not provided in the available records; these must be verified and obtained before journal submission.

8.2 Consent for Publication

Not applicable to the text of this manuscript. Any patient-derived images used in future figures must be fully de-identified before inclusion.

8.3 Availability of Data and Materials

There are no known restrictions; however, the dataset repository, license, and access conditions were not recorded in the available documents. Data will be made available only after the authors confirm ownership of the material, confirm that permission has been granted to use the source, and confirm that de-identification has been completed.

8.4 Funding

No external funding was reported for this final-year project.

8.5 Author Contributions

Mareena Karim, Seena Irshad, Ali Nasir, Imran Maqsood, and Abdullah Ismail contributed to project development, implementation, data-analysis planning, interpretation of results, and preparation of the final report. Sahib Khan supervised the project, provided methodological and technical guidance, and critically reviewed the manuscript. All authors read and approved the final manuscript and are accountable for its accuracy and completeness before submission.

8.6 Conflict of Interest

The authors declare no conflict of interest.

References

O. Boudraa, "Segmentation of 3D dental images using deep learning," Department of Computer Science, XLIM UMR CNRS 7252, University of Limoges, Jul. 21, 2022.

T. T. Nguyen, N. Larrivée, A. Lee, O. Bilaniuk, and R. Durand, "Use of artificial intelligence in dentistry: Current clinical trends and research advances," Journal of the Canadian Dental Association, vol. 87, art. l6, 2021.

A. Ossowska, A. Kusiak, and D. Świetlik, "Artificial intelligence in dentistry—A narrative review," International Journal of Environmental Research and Public Health, vol. 19, no. 6, art. 3449, 2022.

R. Widyaningrum, I. Candradewi, N. R. A. S. Aji, and R. Aulianisa, "Comparison of multi-label U-Net and Mask R-CNN for panoramic radiograph segmentation to detect periodontitis," Imaging Science in Dentistry, vol. 52, no. 3, pp. 227–234, 2022.

L. Boussaid, A. Boukhris, and A. Bensaid, "Teeth U-Net: A segmentation model of dental panoramic X-ray images using hybrid attention mechanisms," Computers in Biology and Medicine, vol. 158, art. 106832, 2023.

Y. Li, Y. Zhang, and Y. Guo, "Automatic teeth segmentation using attention U-Net on dental panoramic X-rays," Journal of Healthcare Engineering, vol. 2022, pp. 1–11, 2022.

M. M. Rahman, M. A. Rahman, and R. Azim, "Denoised encoder-based residual U-Net for precise teeth image segmentation," Scientific Reports, vol. 11, art. 20458, 2021.

S. Kumar, S. Singh, and A. Pandey, "S-R2F2U-Net: A single-stage recurrent residual fusion U-Net for dental image segmentation," Computer Methods and Programs in Biomedicine, vol. 234, art. 107442, 2023.

Frontier in Medical & Health Research

ISSN: (e) 3007-1607 (p) 3007-1593

T. Shan, X. Wang, and X. Huang, "Tooth instance segmentation on panoramic radiographs using U-Net with morphological post-processing," Computers in Biology and Medicine, vol. 144, art. 105387, 2022.

H. J. Chang, K. H. Jeon, and Y. J. Kim, "Instance segmentation and classification of teeth in panoramic radiographs using BB-UNet and YOLOv8," Diagnostics, vol. 12, no. 10, art. 2589, 2022.

A. Madani, M. Moradi, and M. Alizadeh, "FUSegNet: Fusion-based U-Net with attention gates for teeth segmentation and orientation detection," IEEE Access, vol. 10, pp. 70712–70724, 2022.

H. Choi, D. Lee, and S. Park, "Apical lesion segmentation on panoramic radiographs using a U-Net approach," BMC Oral Health, vol. 23, no. 1, pp. 1–10, 2023.

D. Lee, J. Cho, and S. Shin, "Deep learning-based segmentation of dental restorations on panoramic radiographs," Diagnostics, vol. 11, no. 3, art. 547, 2021.

H. Zhang, W. Lu, and T. Yu, "Lightweight deep learning framework for panoramic dental image segmentation," Computers in Biology and Medicine, vol. 136, art. 104699, 2021.

F. Jan, S. Riaz, and T. Mahmood, "Dental Segmentator: Automated CBCT segmentation using nnU-Net," Computers in Biology and Medicine, vol. 141, art. 105124, 2022.

H. Pham, T. Le, and T. Nguyen, "STS-2D/3D-Tooth: A multimodal dental dataset for semi-supervised segmentation," Data in Brief, vol. 47, art. 108979, 2023.

M. Walia and R. Sharma, "Deep learning applications in dental radiographic imaging: A review," Imaging Science Journal, vol. 71, no. 1, pp. 24–38, 2023.

J. Zhao, Y. Yang, and W. Chen, "3D dental image segmentation using CNN with tunicate swarm optimization," Computer Methods and Programs in Biomedicine, vol. 233, art. 107442, 2023.

O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in Proceedings of the Medical Image Computing and Computer-Assisted Intervention Conference, 2015, pp. 234–241.