

Supplemental Materials

Table 1. Quantitative comparison of segmentation performance between Render2Seg and DilatedToothSegNet.

Figure 1. Detection and directional landmark localization results under various dental conditions.

Figure 2. Ablation Study on the Efficacy of GraphCut-based Post-processing for Tooth-wise Segmentation.

Figure 3. Comparison Between Centroid-Based and Anatomical Landmark-Guided Rotated Bounding Boxes for Tooth Patch Extraction.

Table 1. Quantitative comparison of segmentation performance between Render2Seg and DilatedToothSegNet.

Evaluation was performed on the Teeth3DS dataset for both upper and lower jaws. Values are reported as mean $\pm$ standard deviation. Although Render2Seg exhibits slightly higher variance due to occasional detection failures during the bounding box stage, it consistently outperforms the baseline across all metrics.

Method	Jaw	mIoU	IoU	Dice	Accuracy	Precision	Recall	F1 score
DilatedToothSegNet	Upper	0.8541 $\pm$	0.8205 $\pm$	0.8675 $\pm$	0.9316 $\pm$	0.8842 $\pm$	0.8618 $\pm$	0.8675 $\pm$
		0.1185	0.1424	0.1355	0.0692	0.1285	0.1341	0.1355
Render2Seg		0.8873 $\pm$	0.8747 $\pm$	0.9051 $\pm$	0.9401 $\pm$	0.8996 $\pm$	0.9153 $\pm$	0.9051 $\pm$
		0.1164	0.1424	0.1424	0.0758	0.1422	0.1418	0.1424
DilatedToothSegNet	Lower	0.8238 $\pm$	0.7882 $\pm$	0.8395 $\pm$	0.9049 $\pm$	0.8566 $\pm$	0.8327 $\pm$	0.8395 $\pm$
		0.1534	0.1688	0.1645	0.0955	0.1553	0.1643	0.1645
Render2Seg		0.8724 $\pm$	0.8601 $\pm$	0.8901 $\pm$	0.9233 $\pm$	0.8871 $\pm$	0.8961 $\pm$	0.8901 $\pm$
		0.1594	0.1737	0.1778	0.1062	0.1765	0.1785	0.1778

[ref] Lucas Krenmayr, Reinhold von Schwerin, Daniel Schaudt, Pascal Riedel, and Alexander Hafner. Dilatedtoothsegnet: tooth segmentation network on 3d dental meshes through in- creasing receptive vision. Journal of Imaging Informatics in Medicine, 37(4):1846–1862, 2024.

Figure 1. Detection and directional landmark localization results under various dental conditions.

Top row (a–h): Multi-class teeth detection results under various clinical scenarios.

(a–b) show cases with very small teeth or reduced tooth count,
(c–d) represent inclusion or ambush of the 8th tooth (third molar),
(e–f) depict misaligned or rotated teeth, and
(g–h) include cases with one or more missing teeth.

Bounding boxes are color-coded by tooth class, and arrows highlight detection challenges.

Bottom row (i–p): Tooth directional landmark localization results.

(i–j) display results for normal dentition with consistent mesial/distal landmark predictions,

(k–l) correspond to missing tooth cases where landmark prediction remains robust,


(m–p) illustrate crowded dentition cases with close inter-tooth spacing, where the model successfully preserves anatomical alignment.

Blue arrows indicate regions with complex conditions or challenging detection environments.

Multi-Class Teeth Detection result

(a,b) : very small teeth or a small number of teeth,

(c,d) : 8th tooth ambush, (e,f) : misaligned teeth, (g,h) : missing teeth



Tooth Landmark Localization

(i,j) : normal dentition, (k,l) case 2 : missing teeth, (m-p) case 3 : crowded teeth




Figure 2. Ablation Study on the Efficacy of GraphCut-based Post-processing for Tooth-wise Segmentation.

(a) Lower teeth-wise and (b) upper teeth-wise segmentation performance with and without GraphCut-based post-processing. For both plots, dashed lines indicate results without post-processing, and solid lines indicate results with GraphCut refinement. Multiple evaluation metrics (mIoU, IoU, Dice, Accuracy, Precision, Recall, and F1-score) are reported for each tooth number, showing consistent performance improvements achieved by the GraphCut refinement step.





Figure 3. Comparison Between Centroid-Based and Anatomical Landmark-Guided Rotated Bounding Boxes for Tooth Patch Extraction.

(a) Centroid-based bounding box extraction: Tooth patches are cropped using the centroid and a fixed margin, disregarding anatomical orientation. This approach may include unnecessary surrounding tissues and misalign with the tooth's natural axis.

(b) Anatomical landmark-guided rotated bounding box extraction: Bounding boxes are oriented along the distal–mesial axis using clinically relevant landmarks, enabling more anatomically accurate cropping that better captures the crown region while minimizing irrelevant areas.

(a) Centroid-based bounding box extraction



(b) Anatomical landmark-guided rotated bounding box extraction




Figure 4. Impact of Detection Quality and Dentition Orientation on Segmentation Accuracy

This figure illustrates how variations in detection quality and overall dentition orientation can affect the fidelity of tooth-level segmentation results. In cases where tooth detections are incomplete or missing (a), the downstream segmentation stage may omit entire teeth or produce fragmented predictions, as the cropped patches fail to capture the full target region. Similarly, when the entire dentition is rotated relative to the orientation seen during training (b), the extracted bounding boxes may not optimally encompass the tooth surfaces, leading to reduced segmentation precision. These examples highlight the importance of robust detection performance and orientation-invariant handling to maintain segmentation accuracy across diverse clinical presentations.

