



Published in final edited form as:

*Mach Learn Med Imaging*. 2021 September ; 12966: 606–614. doi:10.1007/978-3-030-87589-3\_62.

## SkullEngine: A Multi-Stage CNN Framework for Collaborative CBCT Image Segmentation and Landmark Detection

Qin Liu<sup>1</sup>, Han Deng<sup>2</sup>, Chunfeng Lian<sup>1</sup>, Xiaoyang Chen<sup>1</sup>, Deqiang Xiao<sup>1</sup>, Lei Ma<sup>1</sup>, Xu Chen<sup>1</sup>, Tianshu Kuang<sup>2</sup>, Jaime Gateno<sup>2</sup>, Pew-Thian Yap<sup>1</sup>, James J. Xia<sup>2</sup>

<sup>1</sup>Department of Radiology and Biomedical Research Imaging Center (BRIC), University of North Carolina at Chapel Hill, NC, USA

<sup>2</sup>Department of Oral and Maxillofacial Surgery, Houston Methodist Hospital, TX, USA

### Abstract

Accurate bone segmentation and landmark detection are two essential preparation tasks in computer-aided surgical planning for patients with craniomaxillofacial (CMF) deformities. Surgeons typically have to complete the two tasks manually, spending ~12 hours for each set of CBCT or ~5 hours for CT. To tackle these problems, we propose a multi-stage coarse-to-fine CNN-based framework, called SkullEngine, for high-resolution segmentation and large-scale landmark detection through a collaborative, integrated, and scalable JSD model and three segmentation and landmark detection refinement models. We evaluated our framework on a clinical dataset consisting of 170 CBCT/CT images for the task of segmenting 2 bones (midface and mandible) and detecting 175 clinically common landmarks on bones, teeth, and soft tissues. Experimental results show that SkullEngine significantly improves segmentation quality, especially in regions where the bone is thin. In addition, SkullEngine also efficiently and accurately detect all of the 175 landmarks. Both tasks were completed simultaneously within 3 minutes regardless of CBCT or CT with high segmentation quality. Currently, SkullEngine has been integrated into a clinical workflow to further evaluate its clinical efficiency.

### Keywords

Cone-Beam Computed Tomography (CBCT) Image; Segmentation; Landmark Detection

## 1 Introduction

Accurate bone segmentation and landmark detection are two fundamental tasks in preparing cone-beam computed tomography (CBCT) or computed tomography (CT)<sup>3</sup> scans for use in computer-aided surgical simulation to treat patients with craniomaxillofacial (CMF) deformities. In current clinical practice, it takes at least a day and a half for a surgeon or a trained operator to manually perform both tasks to obtain the CBCT segmentation masks and landmark coordinates, which are time-consuming, labor-intensive, and error-

ptyap@med.unc.edu, jxia@houstonmethodist.org.

<sup>3</sup>For brevity, in the rest of the paper CBCT refers to both CBCT and CT. CBCT is more frequently used clinically.

prone. Therefore, there is an urgent need to develop a reliable automatic segmentation and landmark detection method for clinical use.

Automatic CMF bone segmentation and landmark detection are practically challenging due to the complex CMF anatomy, significant variations in appearance (especially in patients with severe deformities), and large image sizes (up to  $768 \times 768 \times 576$ ). Most existing methods, including conventional methods [3, 12, 14] and CNN-based methods [6, 10, 13, 17], formulate the segmentation and the landmark detection as two independent tasks without considering their inherent relationship (e.g., landmarks usually lie on the boundaries of segmented bone regions).

In recent years, CNN-based joint segmentation and landmark detection (JSD) approaches [7, 16] have been proposed to combine the two tasks via multi-task learning. In [7], the authors proposed a multi-task dynamic transformer network (DTNet) for concurrently segmenting mandible and detecting 64 landmarks. In [16], the authors proposed a context-guided multi-task fully convolutional network for jointly segmenting two bony structures (i.e., midface and mandible) and 15 bony landmarks. However, these approaches have three major drawbacks that hinder them from being integrated into clinical practice: 1) they cannot detect large-scale landmarks (e.g., over 100 landmarks) due to the limited graphics processing unit (GPU) memory, 2) they cannot specifically refine segmentation in regions that are important for surgical planning (e.g., regions with thin bones), 3) they are non-scalable because the segmentation and landmark detection tasks are highly coupled in a single network.

To tackle these issues, we propose a coarse-to-fine CNN-based framework, the SkullEngine, for high-resolution bone segmentation and large-scale landmark detection through a collaborative, integrated, and scalable JSD model and three refinement models. The goal of SkullEngine is to segment 2 bones (i.e., midface and mandible) and 175 landmarks (i.e., 66 for bones, 68 for teeth, and 41 for soft tissues) from a CBCT image. SkullEngine achieves this goal in two sequential stages: coarse and refinement. In the coarse stage, a scalable JSD model, a combination of 3D U-Net-based segmentation and landmark detection models, takes a down-sampled image as input for coarse segmentation and global landmark detection (i.e., all landmarks except the tooth landmarks). The tooth landmarks are not detected in this stage because they are close together in a small region and therefore need to be detected in a higher resolution. In the refinement stage, based on the coarse segmentation mask and global landmarks achieved with the previous stage, region-of-interest volumes are cropped from the original CBCT image for further segmentation refinement and tooth landmark detection.

A major technical contribution of the proposed SkullEngine is that our new JSD model is scalable and modular compared with previous JSDs. In addition, a major clinical contribution of SkullEngine is its clinical effectiveness. The use of SkullEngine can significantly reduce CBCT data preparation time from ~12 hours, or CT from ~5 hours, to 3 minutes for both clinically challenging tasks of high-resolution CBCT segmentation and large-scale landmark detection. We demonstrate in this study the accuracy of SkullEngine using a clinical dataset containing 92 CBCT scans and 78 CT scans of patients with CMF

deformities. We have already integrated SkullEngine into a clinical workflow to continue to evaluate its efficiency in daily clinical practice.

## 2 Methods

SkullEngine is a coarse-to-fine CNN-based framework that consists of two stages: first coarse then refinement (Fig. 1). In Section 2.1, we describe the coarse stage, in which a scalable JSD model is developed for coarse segmentation and global landmark detection (i.e., bony and facial landmarks). In Section 2.2, we describe the refinement stage, in which two segmentation models and one landmark detection model are developed for bone segmentation refinement and local tooth landmark detection, respectively.

### 2.1 Scalable JSD Model for Coarse Segmentation and Global Landmark Detection

Existing JSD models [7, 16] are non-scalable due to their ad-hoc network design for predefined segmentation and landmark detection tasks using different loss functions. With the breakthrough of U-Net [1, 11], as well as its variants [4, 9], it is now possible to use U-Net as a building block to develop a scalable JSD model for unifying the segmentation and landmark detection tasks by taking both of them as a voxel-classification task. In addition, if the segmentation and landmark detection models have the same network structure, we can use transfer learning techniques [2] to train the models more efficiently. Inspired by these ideas, we propose a simple and scalable JSD model (Fig. 2). For segmentation, the ground truth is a voxel-level multi-class mask (Fig. 2 bottom-left). For landmark detection, the ground truth is also a voxel-level landmark mask (Fig. 2 top-left), which are generated from landmark-level coordinates. We generate the landmark mask by assigning voxels to the label of a given landmark (e.g., the  $i$ -th label represents the  $i$ -th landmark) if these voxels belonging to the neighborhood of that landmark. We define the neighborhood of each landmark as a sphere with a predefined radius, which is a hyper-parameter we empirically set to 3 voxels.

**Construction of the Scalable JSD Model**—Our JSD model consists of a segmentation model and two landmark detection models. The segmentation model is used for coarse bone segmentation of the midface and the mandible, whereas the two landmark detection models are used for detecting 66 bony and 41 facial global landmarks, respectively. However, the tooth landmark detection model is not included in the scalable JSD because the resolution of the down-sampled volume is too coarse to achieve any meaningful results.

**Training and Inference**—In the training phase, the input images are down-sampled to a fixed resolution (e.g.,  $2.0 \text{ mm}^3$  in our experiments) for training both the segmentation and landmark detection models. The segmentation model is first trained from scratch. The two landmark detection models are then initialized and fine-tuned using the parameters from the segmentation model. All are trained using the Focal loss function [8]. In the inference phase, the input image is down-sampled to the same resolution as in the training phase. We then run the three models independently for inference.

## 2.2 Bone Segmentation Refinement and Local Landmark Detection

The refinement stage aims to refine the bone segmentation results based on the coarse segmentation results and to detect the tooth landmarks that are undetectable in the coarse stage. As shown in Fig. 1 (right), the original image is cropped based on the coarse mask and global landmarks. Volume (a) in Fig. 1 is a global volume from cropped from the original image that contains the whole skull. The global refinement model uses a patch-based training and inference method for high-resolution segmentation in volume (a) (e.g., by cropping patches with a resolution of  $0.4 \text{ mm}^3$ ). Volume (b) in Fig. 1 shows a volume with thin facial bone. In fact, we need to crop two such volumes (for both the left and right facial bones) from the CBCT image using the paired right and left bony landmarks in the thin bone region as centers. Volume (c) in Fig. 1 is the tooth volume that is cropped also based on the bony landmarks. We train a tooth landmark detection model only in this region with a relative higher resolution (e.g.,  $0.8 \text{ mm}^3$ ). During inference, we first crop the tooth volume that has the same size and resolution as the training patch based on the already detected bony landmarks, and then feed the cropped volume for tooth landmark detection. Finally, the segmentation mask obtained by the global and local refinement models are merged and zero-padded to the original size.

## 3 Experiments and Results

### 3.1 Materials

We evaluated the effectiveness of the proposed framework using 92 CBCT and 78 CT scans, which were randomly selected from our digital archive of 170 patients who had already undergone surgery in treating their CMF deformities (Table 1). The study was approved by Institutional Review Board (Pro00013802). All personal information were de-identified prior to the study. Following clinical standard of care protocol, experienced CMF surgeons used currently available tools to generate segmentation masks and annotate landmark for all 170 scans as ground truth [15]. The segmentation labels included two bony masks: midface and mandible. The 175 landmarks included 66 on both midface and mandible, 68 on both upper and lower teeth, and 41 on facial soft tissues. The average time of creating ground truth was 12 hours for each set of CBCT and 5 hours for each set of CT. Fig. 3 shows a random example of CBCT scan and its ground truth.

We randomly divided the dataset into three groups: 119 scans (70%) for training, 17 (10%) for validation, and 34 (20%) for testing. We applied a stratified sampling strategy to ensure each group included a balanced portion of CBCT and CT scans. The segmentation and landmark detection tasks were completed using our proposed coarse-to-fine SkullEngine framework. Finally, the achieved results were compared with two state-of-the-art competing methods: 3D U-Net [1] and its upgraded variant based on PointRend [5]. We implemented the two methods for both segmentation and landmark detection tasks.

During the evaluation, the computational speed was calculated, starting from the input of CBCT image into SkullEngine until the output of the segmentation masks and landmark coordinates was completed. For segmentation task, we used Dice similarity coefficient (DSC), sensitivity (SEN), and positive prediction value (PPV) to evaluate the segmentation

accuracy using means and standard deviations (SDs). For landmark detection task, we used root mean squared error (RMSE) and true positive rate (TPR) as metrics to evaluate landmark detection accuracy.

### 3.2 Results

For each testing dataset, both segmentation and landmark detection tasks were completed simultaneously within 3 minutes regardless of CBCT or CT. In contrast, a trained operator spent ~12 hours for a set of CBCT and ~5 hours for CT when the current clinical standard method was used. It indicates SkullEngine has a high degree of efficiency.

Table 2 shows the comparison results of segmentation accuracy among the three methods. The results clearly show that our SkullEngine is superior. Fig. 4 shows 2 randomly selected examples. Note that in the thin bone areas are often misidentified as holes due to the low bony density. Since SkullEngine has a refinement model for segmentation, it is capable of solving the “hole” problem in the thin bone regions, unlike the other two competing methods. Table 3 shows the comparison results of landmark detection accuracy between the 3 methods. The results clearly show that our SkullEngine is more accurate. Fig. 5 (left) shows the detection results of SkullEngine on bones, teeth, and facial soft-tissue landmarks, respectively. Fig. 5 (right) shows the error distribution of all 175 landmarks.

### 3.3 Implementation details

We conducted our experiments on a standard workstation equipped with Intel dual-Xeon E5 CPUs, and a single NVidia Titan XP GPU with 12 GB memory. The SkullEngine was implemented and trained using Python 3.7 and Pytorch 1.7.

## 4 Conclusion

In this work, we have proposed a clinical practical framework for high-quality segmentation and large-scale landmark detection from skull CBCT/CT scans. Our multi-stage framework, the SkullEngine, collaboratively integrates segmentation and landmark detection models to maximize the overall performance. The experimental results showed its superior accuracy when compared to the state-of-the-art methods. The results also showed a significant reduction in labor and time spent on CBCT data preparation from 12 hours to less than 3 minutes with high-quality segmentation results, highlighting the practical value of SkullEngine. Currently, we have integrated SkullEngine into our clinical workflow for further evaluating its clinical efficiency.

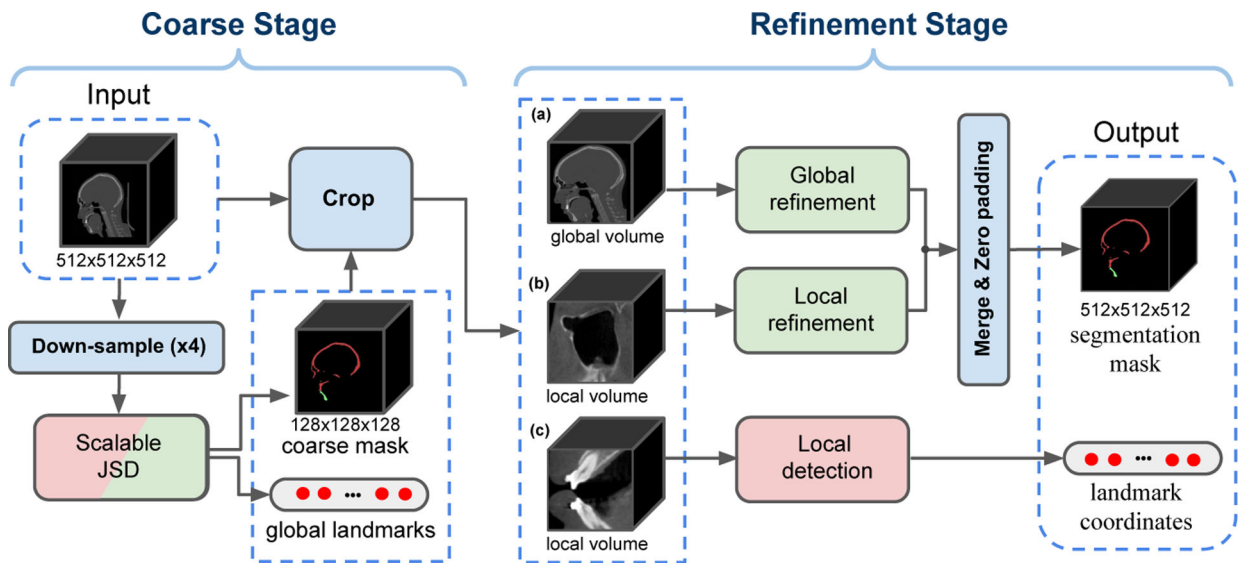
## Acknowledgement

This work was supported in part by United States National Institutes of Health (NIH) grants R01 DE022676, R01 DE027251, and R01 DE021863.

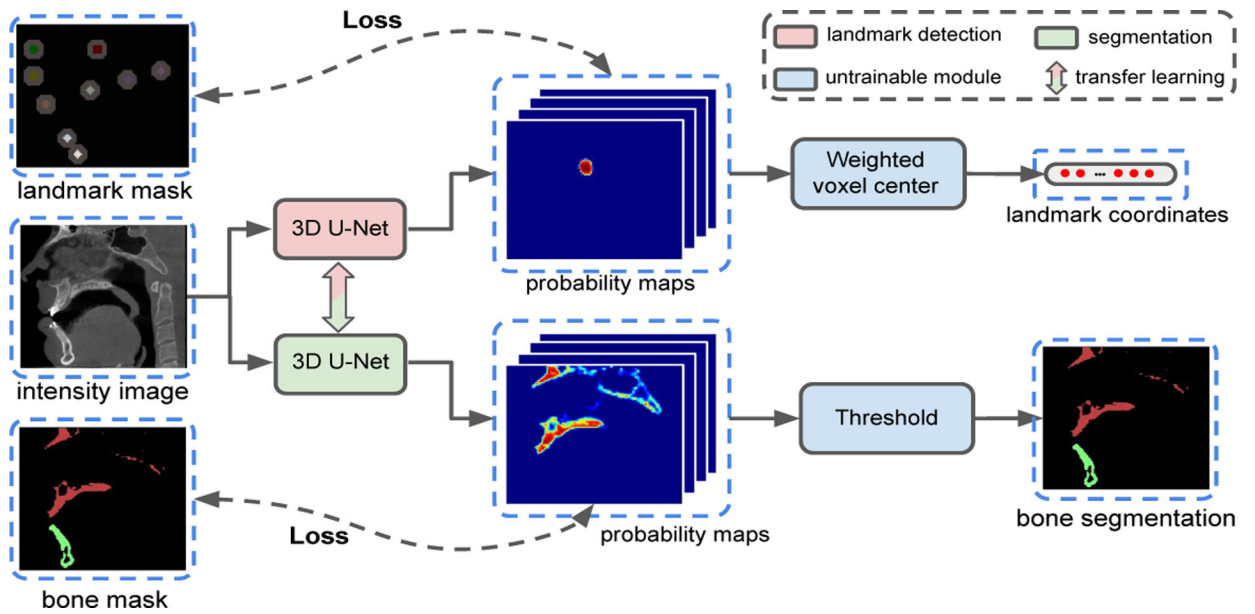
## References

1. Çiçek Ö, Abdulkadir A, Lienkamp SS, Brox T, Ronneberger O: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: International conference on medical image computing and computer-assisted intervention. pp. 424–432. Springer (2016)

2. Dodge J, Ilharco G, Schwartz R, Farhadi A, Hajishirzi H, Smith N: Finetuning pretrained language models: Weight initializations, data orders, and early stopping. arXiv preprint arXiv:2002.06305 (2020)
3. Gupta A, Kharbanda O, Sardana V, Balachandran R: A knowledge-based algorithm for automatic detection of cephalometric landmarks on cbct images. *Int J CARS* 10 (2015)
4. Isensee F, Petersen J, Kohl SAA, Jäger PF, Maier-Hein KH: nnu-net: Breaking the spell on successful medical image segmentation. arXiv preprint arXiv:1904.08128 (2019)
5. Kirillov A, Wu Y, He K, Girshick R: Pointrend: Image segmentation as rendering. arXiv preprint arXiv:2012.08193v2 (2020)
6. Lang Y, Lian C, Xiao D, Deng H, Yuan P, Gateno J, Shen SGF, Alfi DM, Yap PT, Xia JJ, Shen D: Automatic localization of landmarks in craniomaxillofacial cbct images using a local attention-based graph convolution network. In: International conference on medical image computing and computer-assisted intervention. Springer (2020)
7. Lian C, Wang F, Deng HH, Wang L, Xiao D, Kuang T, Lin HY, Gateno J, Shen SGF, Yap PT, Xia JJ, Shen D: Multi-task dynamic transformer network for concurrent bone segmentation and large-scale landmark localization with dental cbct. In: International conference on medical image computing and computer-assisted intervention. Springer (2020)
8. Lin TY, Goyal P, Girshick R, He K, Dollár P: Focal loss for dense object detection. In: IEEE International Conference on Computer Vision (ICCV). pp. 2999–3007. Springer (2017)
9. Milletari F, Navab N, Ahmadi SA: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 2016 Fourth International Conference on 3D Vision (3DV). pp. 565–571. IEEE (2016)
10. Minnema J, van Eijnatten M, Hendriksen AA, Liberton N, Pelt DM, Baten-burg KJ, Forouzanfar T, Wolff J: Segmentation of dental cone-beam ct scans affected by metal artifacts using a mixed-scale dense convolutional neural network. *Medical physics* (2019)
11. Ronneberger O, Fischer P, Brox T: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer Assisted Intervention. pp. 234–241. Springer (2015)
12. Shahidi S, Bahrapour E, Soltanimehr E, Zamani A, Oshagh M, Moattari M, Mehdizadeh A: The accuracy of a designed software for automated localization of craniofacial landmarks on cbct images. *BMC medical imaging* (2014)
13. Torosdagli N, Liberton DK, Verma P, Sincan M, Lee JS, Bagci U: Deep geodesic learning for segmentation and anatomical landmarking. *IEEE Transactions on Medical Imaging* (2018)
14. Wang L, Gao Y, Shi F, Li G, Chen KC, Tang Z, Xia JJ, Shen D: Automated segmentation of dental cbct image with prior-guided sequential random forests. *Medical physics* (2016)
15. Yuan P, Mai H, Li J, Ho DCY, Lai Y, Liu S, Kim D, Xiong Z, Alfi DM, Teichgraber JF, et al. : Design, development and clinical validation of computer-aided surgical simulation system for streamlined orthognathic surgical planning. *International journal of computer assisted radiology and surgery* 12(12), 2129–2143 (2017) [PubMed: 28432489]
16. Zhang J, Liu M, Wang L, Chen S, Yuan P, Li J, Shen SGF, Tang Z, Chen KC, Xia JJ, Shen D: Context-guided fully convolutional networks for joint craniomaxillofacial bone segmentation and landmark digitization. *Medical Image Analysis* (2020)
17. Zheng Y, Liu D, Georgescu B, Nguyen H, Comaniciu D: 3d deep learning for efficient and robust landmark detection in volumetric data. In: International conference on medical image computing and computer-assisted intervention. Springer (2015)

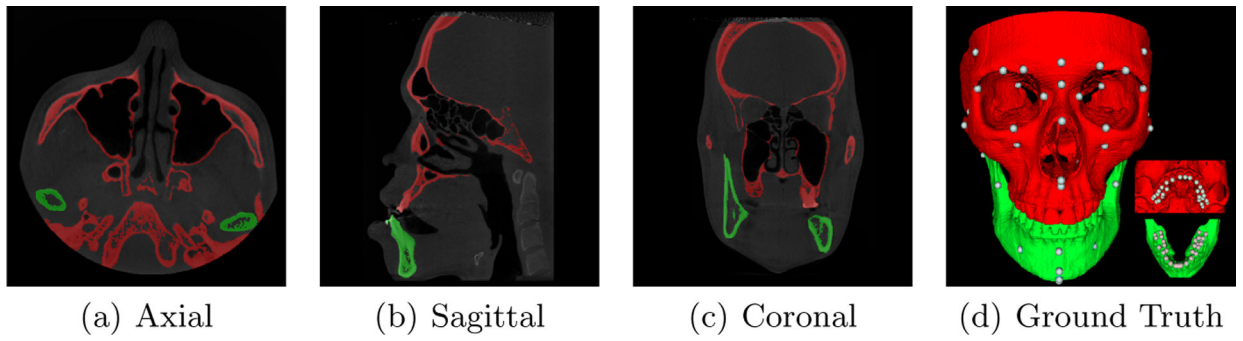


**Fig. 1.** The framework of SkullEngine. The volumes (a) and (b) are cropped for segmentation refinement, while the volume (c) is cropped for tooth landmark detection. The size of each volume on the picture is just for illustration and may be various for different cases.



**Fig. 2.** The training and inference framework for the scalable JSD model. The segmentation model and landmark detection model can be trained as a unified voxel-classification task. In the picture, we only show one segmentation model and one detection model. The transfer learning arrow between the two models means one model can be initialized by the weight from the other model instead of training from scratch.





(a) Axial

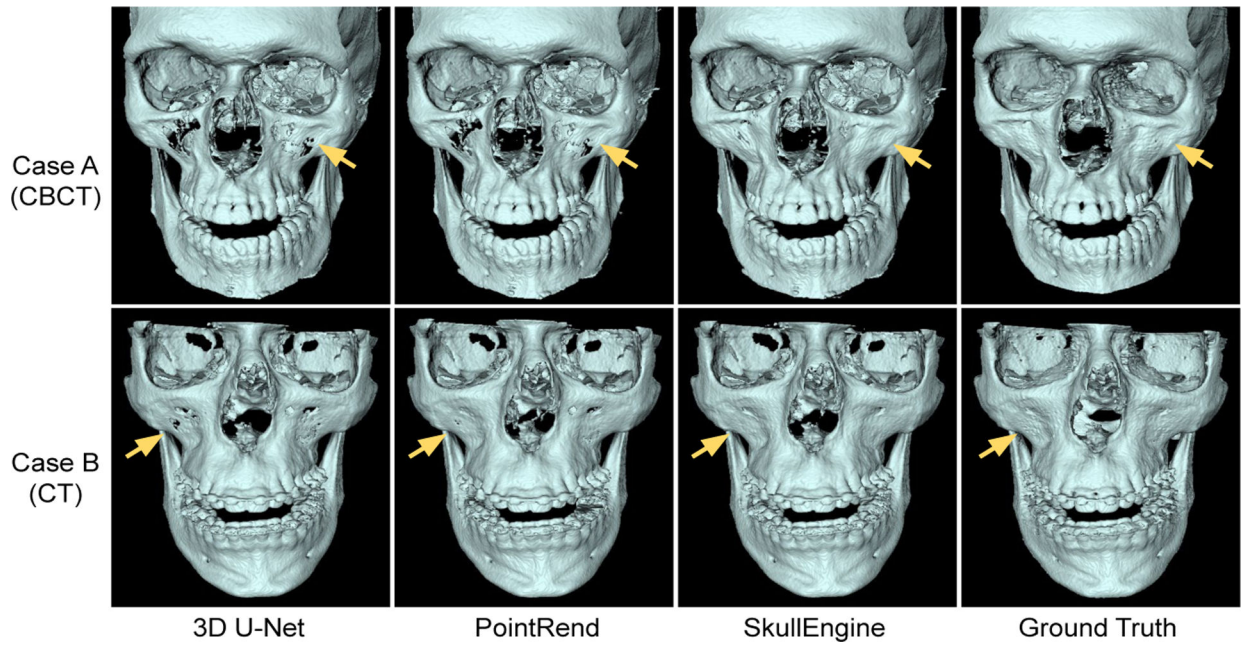
(b) Sagittal

(c) Coronal

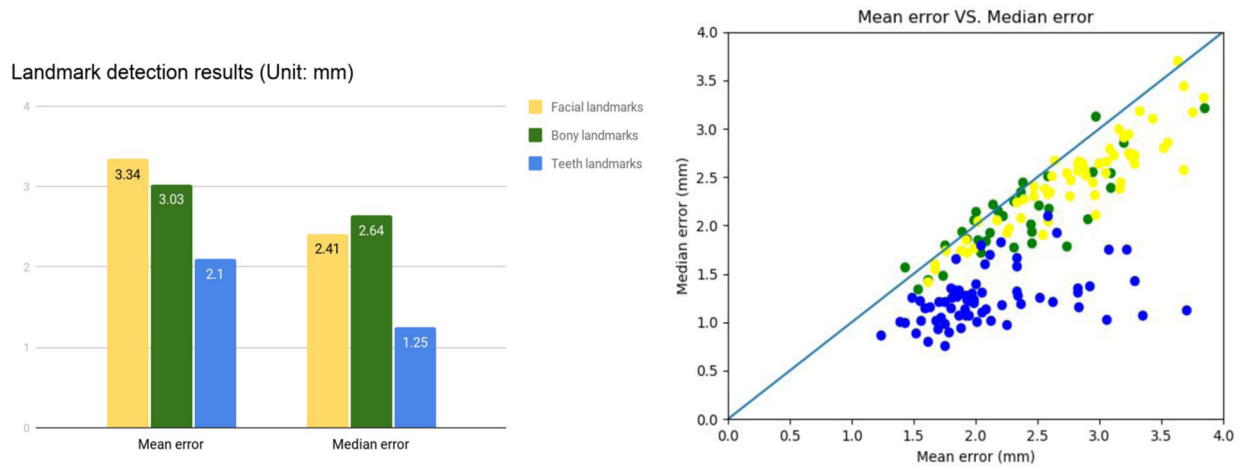
(d) Ground Truth

**Fig. 3.**

An example of CBCT and the ground-truth landmarks. Red label represents the midface, and green represents the mandible. (a) Axial, (b) sagittal, and (c) coronal views, and (d) 3D view of the bony and tooth landmarks marked by spheres.



**Fig. 4.** Comparison of segmentation results for two randomly selected cases. The thin bone areas are pointed by the yellow arrows.



**Fig. 5.** Results of landmark detection of all 175 landmarks using SkullEngine, including 66 bony, 68 teeth, and 41 facial landmarks.

**Table 1.**

Summary of the CBCT/CT dataset.

CBCT/CT Dataset	
Number of scans	170 (CBCT: 92, CT: 78)
Dataset splitting	training 70%, validation 10%, testing 20%
Median spacing (mm <sup>3</sup> )	0.39×0.39×1.0
Median size	512×512×418
Manual segmentations	midface, mandible (both including teeth)
Number of landmarks per scan	175 (bone: 66, teeth: 68, face: 41)

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**Table 2.**

Comparison results of bone segmentation (mean±SD).

Method	Holes on thin bone	Midface			Mandible		
		DSC (%)	SEN (%)	PPV (%)	DSC (%)	SEN (%)	PPV (%)
3D U-Net [1]	Yes	87.9±7.4	83.2±9.5	89.2±7.9	89.4±4.1	91.7±7.1	93.7±4.9
PointRend [5]	Yes	88.3±7.2	84.7±8.9	88.9±8.5	92.4±3.4	91.2±7.4	94.3±5.1
SkullEngine	No	<b>88.5±6.9</b>	<b>85.3±9.3</b>	<b>91.8±7.3</b>	<b>93.5±3.4</b>	<b>92.2±6.1</b>	<b>95.1±4.5</b>

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**Table 3.**

Comparison results of landmark detection (mean±SD).

Method	Bony landmarks		Tooth landmarks		Facial landmarks	
	RMSE	TPR (%)	RMSE	TPR (%)	RMSE	TPR (%)
3D U-Net [1]	3.17±1.79	96.7±3.1	2.58±3.03	97.3±4.1	3.46±3.31	96.4±4.2
PointRend [5]	3.23±2.10	95.0±4.7	2.36±2.96	97.4±3.9	3.28±3.15	97.3±3.7
SkullEngine	<b>3.03±1.96</b>	<b>98.5±2.5</b>	<b>2.10±2.89</b>	<b>98.6±3.5</b>	<b>3.34±3.20</b>	<b>97.5±3.9</b>