

Teeth-SEG: An Efficient Instance Segmentation Framework for Orthodontic Treatment based on Multi-Scale Aggregation and Anthropic Prior Knowledge

Bo Zou *

Tsinghua University
Beijing, China

zoub21@mails.tsinghua.edu.cn

Shaofeng Wang *

Capital Medical University
Beijing, China

2939108747@ccmu.edu.cn

Hao Liu

Tsinghua University
Beijing, China

liuh22@mails.tsinghua.edu.cn

Gaoyue Sun

Imperial College London
London, England

gaoyue.sun22@imperial.ac.uk

Yajie Wang

Tsinghua University, LargeV .Inc
Beijing, China

yj-wang18@mails.tsinghua.edu.cn

FeiFei Zuo

LargeV .Inc
Beijing, China

zuofeifei@largev.com

Chengbin Quan

Tsinghua University
Beijing, China

quancb@tsinghua.edu.cn

Youjian Zhao †

Tsinghua University, Zhongguancun Laboratory
Beijing, China

zhaoyoujian@tsinghua.edu.cn

Abstract

Teeth localization, segmentation, and labeling in 2D images have great potential in modern dentistry to enhance dental diagnostics, treatment planning, and population-based studies on oral health. However, general instance segmentation frameworks are incompetent due to 1) the subtle differences between some teeth' shapes (e.g., maxillary first premolar and second premolar), 2) the teeth's position and shape variation across subjects, and 3) the presence of abnormalities in the dentition (e.g., caries and edentulism). To address these problems, we propose a ViT-based framework named TeethSEG, which consists of stacked Multi-Scale Aggregation (MSA) blocks and an Anthropic Prior Knowledge (APK) layer. Specifically, to compose the two modules, we design a unique permutation-based upscaler to ensure high efficiency while establishing clear segmentation boundaries with multi-head self/cross-gating layers to emphasize particular semantics meanwhile maintaining the divergence between token embeddings. Besides, we collect the first open-sourced intraoral image dataset IO150K, which comprises over 150k intraoral photos, and all photos are annotated by orthodontists using a human-machine hybrid algorithm. Experiments on IO150K demonstrate that our TeethSEG outperforms the state-of-the-art segmentation models on dental image segmentation.

* Equal contribution, † Corresponding author

Our code and dataset will be available at <https://zoub09034.github.io/TeethSEG/>

1. Introduction

Malocclusion, caries, and periodontal disease are the three most common oral cavity diseases, especially the global incidence of malocclusion is a staggering 82.2% [8]. Malocclusion is the misalignment between teeth, jaws, and craniofacial bones caused by genetic or environmental factors during a child's growth. According to the World Dental Federation (FDI), approximately 3.5 billion people worldwide suffer from malocclusion [14], which affects oral health, increasing the risk of caries, periodontal disease, and maxillofacial trauma, also affecting chewing, swallowing, breathing, and pronunciation. Orthodontic treatment is the primary means to cure malocclusion. It utilizes an orthodontic appliance to exert force on teeth in specific directions so that teeth can gradually move and finally achieve the goal of aligning the teeth, reaching the optimal occlusal function, and improving the appearance of the maxillofacial area.

The use of digital technology in oral orthodontics [15, 18, 24–28, 46] has become a popular trend. One of the most widely used applications is integrating artificial intelligence technology to segment oral models and recognize tooth positions automatically. This integration significantly improves the efficiency of treatment plan design and reduces labor costs. Currently, all publicly available intraoral scan data (e.g., [6, 11]) and most teeth segmentation techniques [2, 12, 13, 29, 39] are in 3D space. Although 3D data provides more accurate maxillofacial structure recordings of patients, collecting 3D data is expensive as it requires costly equipment and trained professionals. Furthermore, processing 3D data is challenging and requires high

computing resources, making it unsuitable for large-scale epidemiological screenings and self-inspections. In contrast, obtaining 2D data is relatively simple—a DSLR camera combined with a reflector can obtain standard intraoral dental images. With the improvement in the resolution of mobile phone cameras, individuals can also take their own clear intraoral photos. Dental practitioners can use 2D dental images to document various aspects of a patient’s oral health, such as the alignment, count, color, and general condition of their teeth. By utilizing advanced 2D segmentation algorithms, these images can be analyzed to evaluate tooth crowding, occlusion status, anterior overbite/overjet, and midline alignment of the dental arch.

In recent years, transformer-based models [3–5, 7, 10, 19, 30, 32, 38, 45, 47], have achieved remarkable success in computer vision, quickly dominating various tasks such as image classification, object detection, and semantic segmentation, surpassing traditional CNN models since they can better capture long-range dependencies and unify the modeling of different modalities. However, the teeth segmentation task is distinct from universal semantic segmentations, challenging the state-of-the-art transformer models. Firstly, unlike the apparent differences among object classes in common segmentation tasks, some teeth have similar appearances, such as maxillary first and second premolar. Accurate distinguishing between them requires a complete intraoral assessment. Secondly, due to the varying ages of patients, their teeth are at different stages of development and growth, resulting in different shapes and positions across subjects. Thirdly, caries and tooth loss, common in clinical orthodontic treatment, cause abnormalities in the dentition, which requires models to have strong generalization ability. Finally, to the best of our knowledge, there is no professional annotated 2D teeth segmentation dataset available to support training high-performance models.

To address the current situation, we create the first open-source 2D intraoral scan dataset IO150k, which consists of (1) Challenge80K, 80K rendered images generated from 1,800 3D scans sourced from 3D Teeth Scan Segmentation and Labeling Challenge 2023 [6], (2) Plaster70K, 70K images of 940 oral plaster models made before, during, and after taking the orthodontic treatment, and (3) RGB0.8K, 0.8K RGB standard intraoral photos taken before orthodontic treatment. This dataset has the following key properties: (1) Large: We have collected over 150K images (former dental datasets, e.g., [1, 21, 31, 37], have sizes around 0.1K to 3K) that enable well-trained transformers that are usually more data-hungry than CNN models. (2) Diverse: We cover a wide range of dental malformations (e.g., crowded dentition and edentulism) to ensure the ability to generalize to clinical applications. (3) Professional: The data is annotated by multiple professional orthodontists using a human-machine hybrid algorithm, ensuring accurate tooth position

recognition in complex instances. Please see Appendix A for dataset statistics.

Besides, we propose a novel transformer-based architecture designed against high-performance teeth segmentation named TeethSEG, which has two key components. The first one is Multi-Scale Aggregation Blocks (MSA) that effectively aggregate the visual semantics into trainable class embeddings of each tooth at different scales. The second one is the Anthropropic Prior Knowledge Layer (APK), which imitates the principle of orthodontists to identify teeth, making the segmentation framework more interpretable and perform better, especially when tooth loss happens. Both modules are based on our specially designed multi-head self/cross-gating layers to emphasize valuable components in class embeddings while maintaining the divergence between them. In addition, most dense prediction frameworks [7, 10, 22, 30, 45] use transposed convolution to generate final segmentation maps. Some previous works, like [32], explore transformer-based decoders. However, they have trouble generating clear edges because the embedding sequences’ length is much smaller than the final map size, resulting in mesh-like errors at the segmentation edges. In this paper, we explore replacing upsampling by compressing the intermediate feature dimensions to increase the sequence length, thereby enabling the encoder to learn to store rich local information in different parts of the patch embeddings. Our contributions are summarized as follows:

- We create IO150k, the largest open-source dataset that supports 2D dental segmentation. It covers a wide range of dental malformations and has professional annotations.
- We propose TeethSEG with Multi-Scale Aggregation (MSA) blocks and the Anthropropic Prior Knowledge (APK) layer, and the multi-head cross-gating mechanism and the permutation-based upscaler to form MSA and APK.
- Our experiments demonstrate that TeethSEG outperforms the state-of-the-art general-purpose segmentation models on dental image segmentation.

2. Related Work

Deep learning in Tooth Understanding. Deep learning methods are increasingly used for 3D tooth segmentation. [16, 20, 23, 35, 40, 41, 44]. Mask MCNet [40] combines the Monte Carlo Convolutional Network (MCCNet) with Mask R-CNN to locate each tooth object and segment all the tooth points inside the box. Graph convolutional network-based frameworks (GCN) [33, 34, 42] improve discriminative geometric feature learning for 3D dental model segmentation. TSegNet [13] breaks down dental model segmentation into robust tooth centroid prediction and accurate individual tooth segmentation. DArch [29] proposes to estimate the dental arch and leverage the estimated dental arch to assist the proposal generation of tooth centroids. In summary, the previous method focuses on 3D teeth segmen-

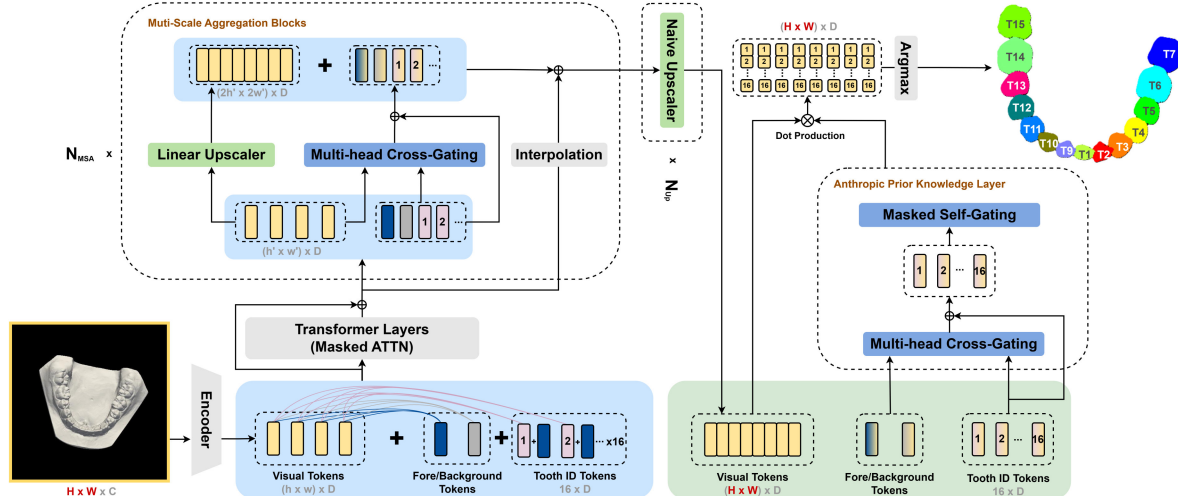


Figure 1. The overview of TeethSEG. We utilize a pretrained encoder to project an intraoral image into a sequence of visual tokens, and a set of trainable class tokens to predict segmentation masks. The multi-scale aggregation (MSA) blocks efficiently aggregate the visual information into class tokens, and the anthropic prior knowledge (APK) layer imposes human judgment into the mask prediction.

tation. We study instance segmentation for 2D intra-oral images, which lowers the data collection and annotation requirement and better supports large-scale epidemiological screenings and self-inspections.

Transformers in Dense Prediction. In recent years, transformers have dominated various tasks. SETR [45] is the first work to adopt ViT as the backbone and develop several CNN decoders for semantic segmentation. Segmenter [32] also extends ViT to semantic segmentation but differs in that it equips a transformer-based decoder. DPT [30] further applies ViT to the monocular depth estimation task via a CNN decoder and yields remarkable improvements. Swin-transformer [22] proposes a shifted-window approach in computing self-attention. BeiT [5] applies masked image modeling as the pretraining tasks to strengthen the encoder. ViT-adapter [10] designs adapter blocks to inject inductive bias for ViTs to enhance performances in dense prediction. These works have achieved remarkable results on general segmentation datasets. However, the teeth segmentation task is distinct from universal semantic segmentations, challenging the state-of-the-art transformer models.

3. Methodology

The primary goal of TeethSEG is to better identify the categories of each individual tooth rather than just distinguishing between tooth areas and background (gingiva) areas. Meanwhile, we make efforts to capture clear segmentation edges with a pure transformer architecture. We choose the multi-model pretrained CLIP encoder as our backbone as its effectiveness has been demonstrated in many downstream tasks. Additionally, its ability to align images with text makes it a strong foundation for expanding TeethSEG into a

multi-modal diagnostic model in the future. In Sec 3.1, we first introduce how to generate segmentation masks based on the pretrained encoder. Then, in Sec 3.2 and Sec 3.3, we introduce the multi-head cross/self-gating mechanism and the permutation-based upscalers (including a naive upscaler and a linear upscaler) that make up our Multi-Scale Aggregation Blocks (MSA) and the Anthropic Prior Knowledge Layer (APK). Finally, in Sec 3.4 and Sec 3.5, we present the details of MSA and APK that are specifically designed for teeth segmentation.

3.1. Overall architecture

An image $X \in \mathbb{R}^{H \times W \times C}$ is encoded into a sequence of visual tokens $x = [x_1, \dots, x_N] \in \mathbb{R}^{N \times D}$ by a pretrained encoder, where $N = hw = HW/P^2$ is the number of visual tokens, (P, P) is the patch size, and D is the dimension of embeddings. Visual tokens x carry rich visual information in the image. We introduce a set of 18 trainable class embeddings to gather the features of the foreground (teeth region), background (gingiva), and each individual tooth. They are divided into foreground/background tokens $CLS_{fb} \in \mathbb{R}^{2 \times D}$ and tooth ID tokens $CLS_{th} \in \mathbb{R}^{16 \times D}$. Before the following computations, as shown in Fig 1, we first add the embedding of the foreground to each tooth ID token in CLS_{th} since all tooth areas should be included in the foreground. After that, We use a M -layer transformer with masked attention to fuse visual tokens and learnable tokens. As shown at the bottom of Fig 1, we apply the attention mask within learnable tokens and only allow visual tokens x to update each token in CLS_{fb} and CLS_{th} at this stage because we want to maximize the dissimilarity within tooth ID tokens CLS_{th} . In this way, we can mitigate the difficulties in distinguishing similar tooth categories. To better merge multi-scale visual semantics into learnable tokens, we utilize MSA blocks in Sec 3.4, which

takes shallow fused x , CLS_{fb} , and CLS_{th} as input to perform deeper feature interaction under different receptive fields. Then, we up-sample the intermediate visual tokens x' to $x' \in \mathbb{R}^{(H \times W) \times D}$ that match the size of the input image X by the permutation-based upscaler in Sec 3.3. Finally, we enable the interactions within learnable tokens under the instruction of human prior knowledge by the APK layer in Sec 3.5. The class masks of each tooth are generated by computing the softmax of the scalar product between x' and tooth ID tokens CLS_{th} as follows:

$$score^{th} = \text{softmax} \left(\frac{x' CLS_{th}^T}{\sqrt{D}} \right), \quad (1)$$

where $score^{th} \in \mathbb{R}^{(H \times W) \times 16}$ is the pixel-wise class score. The \sqrt{D} in the denominator prevents numerical overflow and stabilizes the training. Similarly, the class masks of the foreground and the background are formulated as follows:

$$score^{fb} = \text{softmax} \left(\frac{x' CLS_{fb}^T}{\sqrt{D}} \right). \quad (2)$$

Our model is trained end-to-end with a per-pixel cross-entropy loss consisting of two parts:

$$\mathcal{L}_{th} = -\frac{1}{HW} \sum_{i=1}^{HW} y_i \log(score_i^{th}), \quad (3)$$

$$\mathcal{L}_{fb} = -\frac{1}{HW} \sum_{i=1}^{HW} y_i \log(score_i^{fb}), \quad (4)$$

where y_i is the label of the i -th pixel.

3.2. Multi-Head Cross/Self-Gating Mechanism

We introduce a reusable unit termed cross(self)-gating mechanism for MSA and APK, which takes two arbitrary sub-sequences $V \in \mathbb{R}^{K \times D}$ and $T \in \mathbb{R}^{L \times D}$ as input and performs more efficient feature interactions than commonly used cross-attention after the earlier fusions in the transformer, by exciting or depressing the components of T according to their similarities with V . For a better understanding of the cross-gating Mechanism, we illustrate it and the cross-attention in Fig 2.

There are two key operations of our cross-gating. (1) When K , the length of V , is larger than 1, we sum the similarity matrix $S \in \mathbb{R}^{L \times K}$ over K to form a vector $I \in \mathbb{R}^L$ of importance for token embeddings in T . K functions like the number of multi-heads in the attention mechanism, and every token embedding in V will partially dictate the importance of each token embedding in T . (2) We expand the importance vector I (repeat D times) to match the shape of T . Then, we apply the element-wise product rather than the dot product on the importance matrix and the linear-projected T (*Keys*). The whole process is formulated as follows:

$$Keys = W_k(V), Querys = W_q(T), Values = W_v(T) \quad (5)$$

$$I = \text{repeat} \left(\text{sum} \left(\frac{Querys \cdot Keys}{\|Querys\| \times \|Keys\|} \right) \right), \quad (6)$$

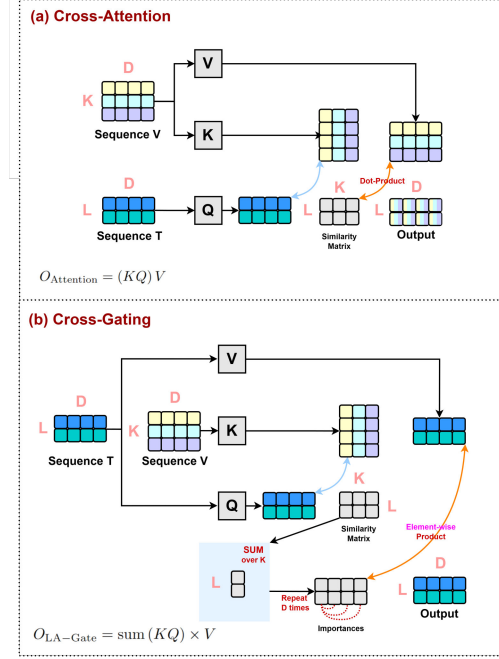


Figure 2. Illustrations of Cross-Attention and our Cross-Gating mechanisms

$$Output = I \odot Values, \quad (7)$$

where \odot denotes element-wise production. In practice, we perform a standard multi-head attention setting [36] on W_k , W_q , W_v , and concatenate outputs of each head.

The most significant characteristic of cross-gating is it can better maintain local diversity within T . In Fig 2 (a), cross-attention’s output displays rows in mixed colors, representing the weighted sum of token embeddings. Consequently, it demonstrates a more global attribute. By contrast, we maintain the uniqueness of colors for cross-gating in Fig 2 (b) because each token embedding is only multiplied by their importance, which is a scalar. This feature is crucial for TeethSEG because the divergence in tooth ID tokens allows us to better distinguish between tooth categories with high similarity. Besides, the interactions brought by commonly used cross-attention (updating embeddings by the weighted sum) can be covered in the previous M -layers transformer when it is applied in MSA and APK since T and V are coming from the same output sequence.

3.3. Permutation-based Upscaler

We upscale the intermediate visual token sequence $x' \in \mathbb{R}^{(h \times w) \times D}$ to $x' \in \mathbb{R}^{(2h \times 2w) \times (D/4)}$ by equally dividing every token embedding with the dimension of \mathbb{R}^D in x' into a sequence with the shape of $\mathbb{R}^{4 \times (D/4)}$ and permute them. We name this process as **naive upscaler** (on the top of Fig 1) since it simply increases the spatial dimensions by compressing the dimension of feature embedding.

However, this naive upscaler performs significantly better than the bilinear interpolation. Applying bilinear interpolation on x' as done in previous techniques [32] does not generate clear segmentation edges. This is because the local information of the interpolated feature map is highly similar, which causes mesh-like errors at the segmentation edge. Besides, the naive upscaler can impose the image encoder to maintain rich local information in its visual tokens.

3.4. Multi-Scale Aggregation Block

The N_{MSA} stacked multi-scale aggregation blocks (MSA) are designed to take in shallow fused x , CLS_{fb} , and CLS_{th} and perform deeper feature interaction under different receptive fields. As shown on the top of Fig 1, each MSA block first uses a cross-gating layer to enhance important semantics in CLS_{th} and CLS_{fb} according to the intermediate visual tokens $x' \in \mathbb{R}^{(h' \times w' \times D)}$, where $h' = 2^k \times h$, $w' = 2^k \times w$, and $k \in [0, N_{\text{MSA}}]$. (In Fig 2, V stands for visual tokens x' , T stands for CLS_{fb} and CLS_{th}). Then, it upsamples x' by what we call **linear upscaler** (on the left of Fig 1), a combination of a linear projection layer W_U and our naive upscaler, where $W_U \in \mathbb{R}^{(D/4 \times D)}$ is used to maintain the embedding dimension D . By stacking N_{MAS} MSA blocks, we can refine the class token embeddings according to multi-scale visual semantics. Additionally, we apply skip connections between each MSA block. The skip connection utilizes the bilinear interpolation to upsample x .

3.5. Anthropic Prior Knowledge Layer

We propose a scalable modular APK to introduce human prior knowledge into the segmentation process. In this paper, we summarize three prior rules based on the annotation experience of orthodontists in complex situations such as tooth loss or abnormal tooth counts in the dental arch. The first rule states that the region requiring tooth ID labeling should not be in the background area, such as the gingiva. The second rule emphasizes the importance of considering the morphological structure of the adjacent teeth on the left and right of each tooth. Finally, the third rule suggests investigating the morphological structure of the contralateral teeth of each tooth based on the symmetry of tooth growth.

To comply with these rules, APK first utilizes a cross-gating layer, which takes CLS_{fb} and CLS_{th} as input (In Fig 2, V represents CLS_{fb} and T represents CLS_{th}), to emphasize the knowledge of foreground in tooth ID tokens. Then, it sends only the processed tooth ID tokens CLS_{th} into a masked self-gating layer, which means V and T in Equation 5 are the same sequence. This layer has an attention mask on the important matrix I in Equation 6, which only enables interaction between adjacent and contralateral teeth, to meet the second and third rules' requirements. The attention mask is visualized in Appendix B.

4. Experiments

4.1. Dataset annotation and processing.

The data annotation, i.e., teeth segmentation and labeling, was performed in collaboration with 4 orthodontists with over 6 years of clinical training experience. The orthodontists were trained in the FDI tooth notation method [17], as well as how to use the annotation software and adhere to the annotation standard. The annotation standard requires each orthodontist to independently annotate all visible deciduous and permanent teeth in each type of intraoral (3D scans of plaster models and 2D RGB photos) data within 7 days. After three weeks, they review all annotations to correct any errors and missed tooth labels. Note, for 2D intra-oral photos, the annotations include teeth with exposed coronal parts or visible residual crowns and roots in the photos but exclude teeth reflected in the reflector for intraoral photography. The detailed process is depicted in Fig 3.

Our approach to reducing labor costs involves a combination of human and machine annotations. To achieve this, we use FusionAnalyser [43], a dental model analysis tool, for 3D scans (on the bottom of Fig 3). In this process, orthodontists draw the boundary line and identify the corresponding tooth ID for each tooth region. The software then automatically generates 3D segmentation for each tooth. For 2D photos (on the top of Fig 3), we first ask orthodontists to label the central point of each tooth class. We then use SAM [19], an open-source image segmentation framework, to generate segmentation masks based on the human-labeled tooth centers. Finally, we ask orthodontists to verify all auto-generated segmentations for both data types. This process ensures the accuracy of the final segmentation. Finally, we rotate the 3D models and project them onto 2D images with labels in various angles. Our method significantly increases the dataset's sample richness while minimizing sample collection costs and annotation costs. Besides, training on a multi-angle plaster cast also helps to improve the model's tolerance for low-quality intra-oral shots (camera angle skew).

4.2. Experimental Setup

Competing Methods. We compare our approach with the state-of-the-art methods (i.e., DeepLabV3 [9], Segformer [32], Segformer [38], Swin-transformer [22], BeiT [5], and ViT-adapter [10]) of 2D instance segmentation. DeepLabV3 is a powerful DilatedFCN-based model with atrous spatial pyramid pooling introducing rich multi-scale information. Segformer, which uses a masked transformer to generate segmentation masks, is an earlier attempt that brings the vision transformer (ViT) into the field of semantic segmentation. Segformer comes up with using overlapped image patches to increase local continuity for ViT-based models and uses deep-wise convolutions to replace the po-

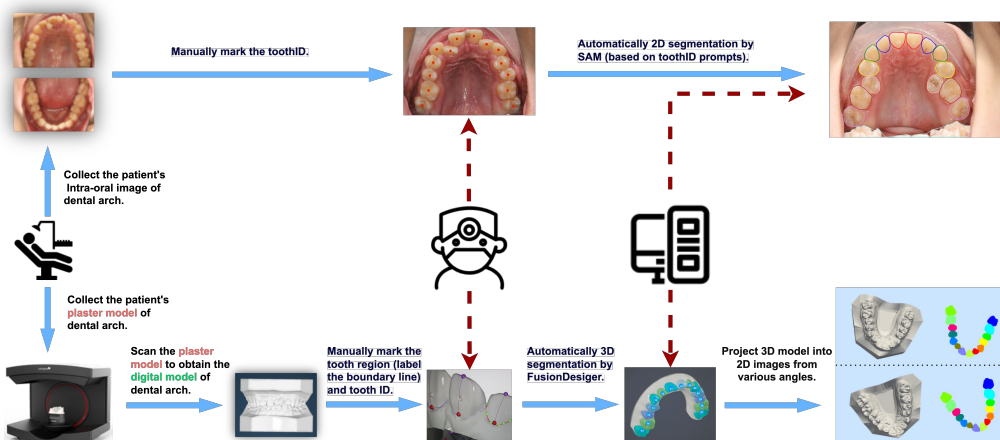


Figure 3. Illustration of our human-machine hybrid data annotation process.

Method	Epoch	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10	T11	T12	T13	T14	T15	T16	mIoU
DeepLab-v3	30	0.85	0.85	0.85	0.89	0.91	0.96	0.87	NaN	0.85	0.83	0.89	0.91	0.83	0.85	0.82	NaN	0.87
Segformer	30	0.86	0.85	0.84	0.86	0.91	0.93	0.87	NaN	0.87	0.84	0.88	0.91	0.77	0.77	0.72	NaN	0.85
Segmenter	30	0.78	0.77	0.74	0.68	0.64	0.77	0.75	NaN	0.78	0.76	0.81	0.86	0.71	0.75	0.69	NaN	0.75
Swin-L	30	0.81	0.79	0.77	0.78	0.82	0.89	0.85	NaN	0.82	0.80	0.83	0.87	0.74	0.75	0.70	NaN	0.80
SwinV2-G	30	0.81	0.79	0.77	0.84	0.84	0.90	0.85	NaN	0.81	0.79	0.83	0.87	0.74	0.76	0.71	NaN	0.80
BeiT-B	30	0.79	0.76	0.76	0.78	0.84	0.90	0.84	NaN	0.80	0.77	0.81	0.85	0.77	0.80	0.74	NaN	0.80
ViT-Adapter-L	30	0.87	0.87	0.84	0.86	0.90	0.93	0.87	NaN	0.88	0.87	0.89	0.87	0.82	0.87	0.82	NaN	0.87
TeethSEG	5	0.92	0.92	0.84	0.89	0.94	0.96	0.89	NaN	0.92	0.92	0.93	0.95	x0.85	0.87	0.82	NaN	0.91

Table 1. Results (mIoU) compared with SOTA methods on the IO150K independent and identically distributed (i.i.d.) test splits

sitional embedding. Swin-transformer proposes a shifted-window approach in computing self-attention, increasing token embeddings’ scale while reducing overhead. BeiT applies masked image modeling (MIM), a token-level autoregression, as the pretraining tasks to strengthen the encoder. ViT-adapter designs adapter block to inject inductive bias for ViTs to enhance performances in dense prediction.

Implementation Details. Our IO150K contains three parts: (1) Challenge80K, 80K rendered images from 3D scans provided by 3D Teeth Scan Segmentation and Labeling Challenge 2023 [6], (2) Plaster70K, 70K images of oral plaster models, and (3) RGB0.8K, 0.8K RGB standard intra-oral photos. Each part has been individually divided into training, validation, and testing splits (please see Appendix A for details). Although all three parts of IO150K support separate training and testing for future studies, in this paper, we first pretrain models on the training split of Challenge80K and test on Challenge80K and Plaster70K testing splits (denoted as **i.i.d. test** (independent and identically distributed), **o.o.d. test** (out of the distribution)), then finetune and test models on RGB0.8K (denoted as **RGB test**). This is because the data in Challenge80K is general tooth data that matches the real-world distribution, while the Plaster80K and RGB0.8K we collected are from patients who accept the orthodontic examination. Compared with Challenge80K, their samples have more abnormalities and

higher complexity. We hope the model trained on general data can adapt to the needs of orthodontic diagnosis (o.o.d. test), and the models trained on a large number of 3D model projections can be transferred to the use of analyzing RGB intra-oral photos for early screening (RGB test).

We report the results of TeethSEG with pretrained CLIP-L/14@336 as the encoder. The embedding dim $D = 768$, the number of layers in the masked transformer $M = 3$, the number of stacked MSA blocks $N_{MSA} = 3$, and the number of stacked naive upscalers $N_{up} = 2$.

4.3. Comparison with Competing Methods

I.I.D. Test Results. The overall detection and segmentation results are presented in Table 1, and we compare these competing methods in the Intersection over Union (IoU) of each tooth class (denoted as T1 to T16. The pre-defined order of classes is shown on the top of Fig 1. NaN means the corresponding class is not shown in our test split.) and the average over all classes (mIoU). The table shows that our TeethSEG achieves the best segmentation performance in each tooth class and improves overall segmentation performance by 4% compared with the state-of-the-art methods. From the table, we can see that for the i.i.d. test, competing methods have similar performance on teeth with a large area (i.e., T6, T7, T14, T15) and have significant performance differences on smaller teeth (i.e., T1, T2, T9, T10). In par-

Method	Epoch	mIoU	
		o.o.d.	RGB
DeepLab-v3	30	0.80	0.80
Segformer	30	0.81	0.69
Segmenter	30	0.68	0.55
Swin-L	30	0.64	0.48
SwinV2-G	30	0.60	0.46
BeiT-B	30	0.78	0.47
ViT-Adapter-L	30	0.79	0.85
TeethSEG	5	0.84	0.91

Table 2. Tooth segmentation results (mIoU) on the IO150K out-of-the-distribution (o.o.d.) test splits and RGB test split. Please see the Appendix for the IoU of each tooth ID.

ticular, methods targeting capturing multi-scale objects (i.e., DeepLab-v3, Segformer, ViT-Adapter) have better performances on smaller teeth. Our TeethSEG uses MSA Blocks to capture multi-scale information and improves 4% to 5% IoU performance on T1, T2, T9, and T10.

O.O.D. Test Results. We compare the o.o.d. performance of TeethSEG with competing methods in Table 2 and find that our method outperforms the state-of-the-art methods with 3% mIoU, which shows TeethSEG’s generalize ability on data ad-hoc to orthodontic treatment. We also visualize randomly picked 4 segmentation predictions of some methods and highlight the incorrect parts in Table 3. From this table, we can find competing methods have different degrees of errors in dealing with complex situations such as missing teeth or irregular tooth arrangements. Interestingly, the methods that perform well in the i.i.d. test still obtain clear tooth segmentation boundaries in this test (i.e., correctly distinguishing between tooth areas and background areas). However, they all have the problem of incorrectly categorizing some teeth as belonging to other tooth categories (in the ground truth, we assigned a unique color to each tooth ID). By contrast, our framework can better identify tooth IDs by incorporating relevant dental arch information, with the help of human prior knowledge via the APK layer in TeethSEG. Please see Appendix C for the full visualization.



Figure 4. Examples of TeethSEG’s segmentation results on IO150K RGB test split.

RGB Test Results. We finetune the models on our RGB0.8K to show the pretrained knowledge on plaster models can be transferred into the RGB domain. Table 2 reports the performances of TeethSEG with competing methods. It shows that TeethSEG brings 6% performance boost, which demonstrates the generalization ability of our frame-

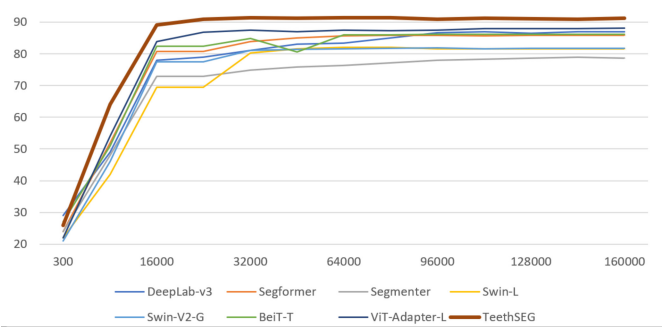


Figure 5. The trend of mIoU changes during the training process.

work. Fig 4 visualizes several randomly picked prediction results of TeethSEG on intra-oral photos from patients before receiving orthodontic treatment. We find that even in cases of obvious dental arch abnormalities, TeethSEG can still accurately segment the tooth area and identify the correct tooth ID. Please see Appendix C for the visual comparison with other methods.

Comparison on Training Speed. Due to the specialized design introduced by TeethSEG for tooth segmentation, its training speed is higher than competing methods. In Fig 5, we visualize the change of mIoU on Challenge80K during the pretraining for all methods.

4.4. Ablation



Figure 6. Comparison of Bilinear Interpolation to Permutation-based Upscaler. (left) Ground Truth, (middle) Bilinear Interpolation, (right) Permutation-based Upscaler.

Permutation-based Upscaler vs. Bilinear Interpolation.

Previous transformer-only decoders used bilinear interpolation for scaling the intermediate feature map to match the size of the output, causing errors at segmentation edges. Besides, the interpolated enlarged image does not introduce new information to local areas, which prohibits the model from learning multi-scale semantics during training. In Fig 6, we visualize the background segmentation result generated by using CLS_{fb} in Sec. 3.1 and compare it with the result from a variant of replacing all linear upscales and naive upscales with bilinear interpolation. We also quantify the performance difference in Table 4.

Effectiveness of each component. To verify the effectiveness of our permutation-based upscales, cross-gating mechanism, MSA blocks, and the APK layer, we design six variants and report their performances in Table 4. (a) We re-

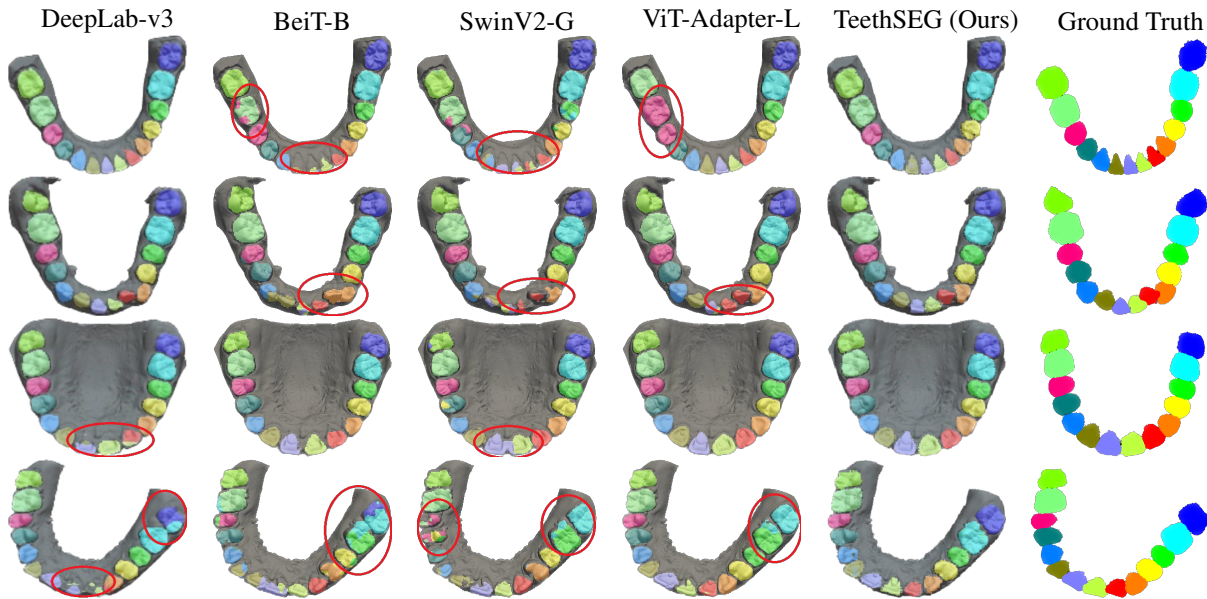


Table 3. The visual comparison of segmentation results (o.o.d test), as well as the corresponding ground truth.

Method	Module						i.i.d. test	o.o.d. test
	Permute-UP	Bilinear-UP	Cross-Gate	Cross-ATT	MSA	APK		
(a)	✗	✓	✓	✗	✓	✓	0.72	0.67
(b)	✓	✗	✗	✓	✓	✓	0.89	0.80
(c)	✓	✗	✓	✗	✗	✓	0.87	0.73
(d)	✓	✗	✓	✗	✓	✗	0.89	0.76
(e)	✓	✓	-	-	✗	✗	0.79	0.70
(f)	✓	✗	✓	✗	✓	✓	0.91	0.84

Table 4. Analysis of the effectiveness of each module.

place all linear upscalers and naive upscalers with bilinear interpolation. (b) We replace our proposed cross/self-gating with the standard cross/self-attention in the MSA blocks and the APK layer. In (c), (d), and (e), we explore the influence of removing the MSA blocks or the APK layer, as well as removing them all. (f) is our best variant with all specially designed modules.

More Ablations. We validate the influence of the scale of the image encoder, the resolution of the input images, and the reasonable choice of hyper-parameters in Appendix D.

5. Conclusion

In this paper, we study the 2D image segmentation. To address the gap in research in this field, we create an open-source dataset called IO150k, which covers a wide range of dental malformations and is intended to serve as a resource for future research. Furthermore, starting from the particularity of the dental segmentation, we design TeethSEG, which surpasses the performance of the state-of-the-art segmentation models. This model includes two modules: Multi-Scale Aggregation (MSA) block and Anthropoc Prior

Knowledge (APK) layer. The former effectively aggregates the visual semantics into class embeddings at different scales, and the latter imitates the principle of orthodontists to identify teeth. To realize MSA and APK, we developed a cross/self-gating mechanism for efficient deep feature interaction, as well as a permutation-based upscaler to generate clear segmentation edges and maintain local information in image patch embeddings. Experiments conducted in this paper demonstrate the effectiveness of our model and indicate that pretraining on plaster models can facilitate the segmentation of intra-oral images, which has the potential to assist large-scale epidemiological screenings and self-inspections.

Acknowledge

This work is supported in part by the Beijing Natural Science Foundation (No. L222024), the National Natural Science Foundation of China (No. 62394322), and the Beijing Hospitals Authority Clinical medicine Development of special funding support (No. ZLRK202330).

References

- [1] Aqsa Ajaz and D Kathirvelu. Dental biometrics: Computer aided human identification system using the dental panoramic radiographs. In *2013 international conference on communication and signal processing*, pages 717–721. IEEE, 2013. **2**
- [2] Ammar Alsheghri, Farnoosh Ghadiri, Ying Zhang, Olivier Lessard, Julia Keren, Farida Cheriet, and François Guibault. Semi-supervised segmentation of tooth from 3d scanned dental arches. In *Medical Imaging 2022: Image Processing*, pages 766–771. SPIE, 2022. **1**
- [3] Anonymous. LLaMA-ecitor: General instruction tuning via indirect feature interaction. In *Conference on Computer Vision and Pattern Recognition 2024*, 2024. **2**
- [4] Anonymous. Videodistill: Language-aware vision distillation for video question answering. In *Conference on Computer Vision and Pattern Recognition 2024*, 2024.
- [5] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. **2, 3, 5**
- [6] Achraf Ben-Hamadou, Oussama Smaoui, Ahmed Rezik, Sergi Pujades, Edmond Boyer, Hoyeon Lim, Minchang Kim, Minkyung Lee, Minyoung Chung, Yeong-Gil Shin, et al. 3dteethseg’22: 3d teeth scan segmentation and labeling challenge. *arXiv preprint arXiv:2305.18277*, 2023. **1, 2, 6**
- [7] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. **2**
- [8] Niccoló Cenzato, Anna Nobili, and Cinzia Maspero. Prevalence of dental malocclusions in different geographical areas: scoping review. *Dentistry Journal*, 9(10):117, 2021. **1**
- [9] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. **5**
- [10] Zhe Chen, Yuchen Duan, Wenhai Wang, Junjun He, Tong Lu, Jifeng Dai, and Yu Qiao. Vision transformer adapter for dense predictions. *arXiv preprint arXiv:2205.08534*, 2022. **2, 3, 5**
- [11] Weiwei Cui, Yaqi Wang, Yilong Li, Dan Song, Xingyong Zuo, Jiaojiao Wang, Yifan Zhang, Huiyu Zhou, Bung san Chong, Liaoyuan Zeng, et al. Ctooth+: A large-scale dental cone beam computed tomography dataset and benchmark for tooth volume segmentation. In *MICCAI Workshop on Data Augmentation, Labelling, and Imperfections*, pages 64–73. Springer, 2022. **1**
- [12] Weiwei Cui, Yaqi Wang, Qianni Zhang, Huiyu Zhou, Dan Song, Xingyong Zuo, Gangyong Jia, and Liaoyuan Zeng. Ctooth: a fully annotated 3d dataset and benchmark for tooth volume segmentation on cone beam computed tomography images. In *International Conference on Intelligent Robotics and Applications*, pages 191–200. Springer, 2022. **1**
- [13] Zhiming Cui, Changjian Li, Nenglu Chen, Guodong Wei, Runnan Chen, Yuanfeng Zhou, Dinggang Shen, and Wenping Wang. Tsegnet: An efficient and accurate tooth segmentation network on 3d dental model. *Medical Image Analysis*, 69:101949, 2021. **1, 2**
- [14] FDI World Dental Federation. Malocclusion in orthodontics and oral health: Adopted by the general assembly: September 2019, san francisco, united states of america. *International Dental Journal*, 70(1):11, 2020. **1**
- [15] Hye-Won Hwang, Ji-Hoon Park, Jun-Ho Moon, Youngsung Yu, Hansuk Kim, Soo-Bok Her, Girish Srinivasan, Mohammed Noori A Aljanabi, Richard E Donatelli, and Shin-Jae Lee. Automated identification of cephalometric landmarks: Part 2-might it be better than human? *The Angle Orthodontist*, 90(1):69–76, 2020. **1**
- [16] Joon Im, Ju-Yeong Kim, Hyung-Seog Yu, Kee-Joon Lee, Sung-Hwan Choi, Ji-Hoi Kim, Hee-Kap Ahn, and Jung-Yul Cha. Accuracy and efficiency of automatic tooth segmentation in digital dental models using deep learning. *Scientific reports*, 12(1):9429, 2022. **2**
- [17] S Keiser-Nielsen. Fédération dentaire internationale two-digit system of designating teeth. *Int Dent J*, 21:104–106, 1971. **5**
- [18] Jaerong Kim, Inhwan Kim, Yoon-Ji Kim, Minji Kim, Jin-Hyoung Cho, Mihee Hong, Kyung-Hwa Kang, Sung-Hoon Lim, Su-Jung Kim, Young Ho Kim, et al. Accuracy of automated identification of lateral cephalometric landmarks using cascade convolutional neural networks on lateral cephalograms from nationwide multi-centres. *Orthodontics & Craniofacial Research*, 24:59–67, 2021. **1**
- [19] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. **2, 5**
- [20] André Ferreira Leite, Adriaan Van Gerven, Holger Willems, Thomas Beznik, Pierre Lahoud, Hugo Gaëta-Araujo, Myrthel Vranckx, and Reinhilde Jacobs. Artificial intelligence-driven novel tool for tooth detection and segmentation on panoramic radiographs. *Clinical oral investigations*, 25:2257–2267, 2021. **2**
- [21] Yunxiang Li, Guodong Zeng, Yifan Zhang, Jun Wang, Qun Jin, Lingling Sun, Qianni Zhang, Qisi Lian, Guiping Qian, Neng Xia, et al. Agmb-transformer: Anatomy-guided multi-branch transformer network for automated evaluation of root canal therapy. *IEEE Journal of Biomedical and Health Informatics*, 26(4):1684–1695, 2021. **2**
- [22] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. **2, 3, 5**
- [23] Zuozhu Liu, Xiaoxuan He, Hualiang Wang, Huimin Xiong, Yan Zhang, Gaoang Wang, Jin Hao, Yang Feng, Fudong Zhu, and Haoji Hu. Hierarchical self-supervised learning for 3d tooth segmentation in intra-oral mesh scans. *IEEE Transactions on Medical Imaging*, 42(2):467–480, 2022. **2**
- [24] Jun-Ho Moon, Hye-Won Hwang, Youngsung Yu, Min-Gyu Kim, Richard E Donatelli, and Shin-Jae Lee. How much deep learning is enough for automatic identification to be reliable? a cephalometric example. *The Angle Orthodontist*, 90(6):823–830, 2020. **1**

- [25] Soh Nishimoto, Yohei Sotsuka, Kenichiro Kawai, Hisako Ishise, and Masao Kakibuchi. Personal computer-based cephalometric landmark detection with deep learning, using cephalograms on the internet. *Journal of Craniofacial Surgery*, 30(1):91–95, 2019.
- [26] Ji-Hoon Park, Hye-Won Hwang, Jun-Ho Moon, Young-sung Yu, Hansuk Kim, Soo-Bok Her, Girish Srinivasan, Mohammed Noori A Aljanabi, Richard E Donatelli, and Shin-Jae Lee. Automated identification of cephalometric landmarks: Part 1—comparisons between the latest deep-learning methods yolov3 and ssd. *The Angle Orthodontist*, 89(6):903–909, 2019.
- [27] Christian Payer, Darko Štern, Horst Bischof, and Martin Urschler. Integrating spatial configuration into heatmap regression based cnns for landmark localization. *Medical image analysis*, 54:207–219, 2019.
- [28] William R Proffit, Henry W Fields, Brent Larson, and David M Sarver. *Contemporary orthodontics-e-book*. Elsevier Health Sciences, 2018. 1
- [29] Liangdong Qiu, Chongjie Ye, Pei Chen, Yunbi Liu, Xiaoguang Han, and Shuguang Cui. Darch: Dental arch prior-assisted 3d tooth instance segmentation with weak annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20752–20761, 2022. 1, 2
- [30] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12179–12188, 2021. 2, 3
- [31] Gil Silva, Luciano Oliveira, and Matheus Pithon. Automatic segmenting teeth in x-ray images: Trends, a novel data set, benchmarking and future perspectives. *Expert Systems with Applications*, 107:15–31, 2018. 2
- [32] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7262–7272, 2021. 2, 3, 5
- [33] Diya Sun, Yuru Pei, Peixin Li, Guangying Song, Yuke Guo, Hongbin Zha, and Tianmin Xu. Automatic tooth segmentation and dense correspondence of 3d dental model. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 703–712. Springer, 2020. 2
- [34] Diya Sun, Yuru Pei, Guangying Song, Yuke Guo, Gengyu Ma, Tianmin Xu, and Hongbin Zha. Tooth segmentation and labeling from digital dental casts. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pages 669–673. IEEE, 2020. 2
- [35] Sukun Tian, Ning Dai, Bei Zhang, Fulai Yuan, Qing Yu, and Xiaosheng Cheng. Automatic classification and segmentation of teeth on 3d dental model using hierarchical deep learning networks. *IEEE Access*, 7:84817–84828, 2019. 2
- [36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 4
- [37] Ching-Wei Wang, Cheng-Ta Huang, Jia-Hong Lee, Chung-Hsing Li, Sheng-Wei Chang, Ming-Jhih Siao, Tat-Ming Lai, Bulat Ibragimov, Tomaž Vrtovec, Olaf Ronneberger, et al. A benchmark for comparison of dental radiography analysis algorithms. *Medical image analysis*, 31:63–76, 2016. 2
- [38] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34:12077–12090, 2021. 2, 5
- [39] Xiaojie Xu, Chang Liu, and Youyi Zheng. 3d tooth segmentation and labeling using deep convolutional neural networks. *IEEE transactions on visualization and computer graphics*, 25(7):2336–2348, 2018. 1
- [40] Farhad Ghazvinian Zanjani, David Anssari Moin, Frank Claessen, Teo Cherici, Sarah Parinussa, Arash Pourtaherian, Svitlana Zinger, et al. Mask-mcnet: Instance segmentation in 3d point cloud of intra-oral scans. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 128–136. Springer, 2019. 2
- [41] Jianda Zhang, Chunpeng Li, Qiang Song, Lin Gao, and Yukun Lai. Automatic 3d tooth segmentation using convolutional neural networks in harmonic parameter space. *Graphical Models*, 109:101071, 2020. 2
- [42] Lingming Zhang, Yue Zhao, Deyu Meng, Zhiming Cui, Chenqiang Gao, Xinbo Gao, Chunfeng Lian, and Dinggang Shen. Tsgnet: Discriminative geometric feature learning with two-stream graph convolutional network for 3d dental model segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6699–6708, 2021. 2
- [43] Shaofeng Zhang. Fusionanalyser: A novel measurement method and software tool for dental model analysis in orthodontics. *Proceedings of Measurement Science and Technology*, 2023. 5
- [44] Yue Zhao, Lingming Zhang, Yang Liu, Deyu Meng, Zhiming Cui, Chenqiang Gao, Xinbo Gao, Chunfeng Lian, and Dinggang Shen. Two-stream graph convolutional network for intra-oral scanner image segmentation. *IEEE Transactions on Medical Imaging*, 41(4):826–835, 2021. 2
- [45] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6881–6890, 2021. 2, 3
- [46] Zhushi Zhong, Jie Li, Zhenxi Zhang, Zhicheng Jiao, and Xinbo Gao. An attention-guided deep regression model for landmark detection in cephalograms. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part VI 22*, pages 540–548. Springer, 2019. 1
- [47] Bo Zou, Chao Yang, Chengbin Quan, and Youjian Zhao. Spaceclip: A vision-language pretraining framework with spatial reconstruction on text. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 519–528, 2023. 2