









# OralSAM: One-shot Segmentation for Intraoral Ultrasound Videos with Adaptive Feature Correlation and Self-prompting Strategy

Logiraj Kumaralingam<sup>1\*</sup> , Anparasy Sivaanpu<sup>1\*</sup> , Manh-Hai Hoang<sup>1</sup> ,  
Javaneh Alavi<sup>1</sup>, Kim-Cuong T. Nguyen<sup>1</sup> , Kumaradevan Punithakumar<sup>1</sup> ,  
Edmond H. M. Lou<sup>2,3</sup> , Paul Major<sup>4</sup>, and Lawrence H. Le<sup>1,2,4,5</sup>  

<sup>1</sup>Department of Radiology and Diagnostic Imaging, University of Alberta, Edmonton, Alberta, Canada

<sup>2</sup>Department of Biomedical Engineering, University of Alberta, Edmonton, Alberta, Canada

<sup>3</sup>Department of Electrical Engineering, University of Alberta, Edmonton, Alberta, Canada

<sup>4</sup>School of Dentistry, University of Alberta, Edmonton, Alberta, Canada

<sup>5</sup>Department of Physics, University of Alberta, Edmonton, Alberta, Canada  
lawrence.le@ualberta.ca

**Abstract.** Periodontal disease is a leading cause of tooth loss and is linked to systemic conditions such as endocarditis, diabetes, cardiovascular disease, and osteoporosis. Intraoral ultrasound (IUS) videos offer a non-invasive means for diagnosing periodontal structures, but existing segmentation methods rely on extensive manual annotations. We propose OralSAM, a one-shot video segmentation network inspired by the Segment Anything Model (SAM), which requires annotation from only a single frame. Our network integrates an adaptive feature correlation module to capture temporal dependencies and refine segmentation consistency across frames. Additionally, we introduce a self-prompting strategy based on optical flow, dynamically adjusting point prompts based on motion cues in consecutive frames to improve segmentation accuracy. To further enhance robustness, we incorporate a self-correction mechanism that refines mask embeddings adaptively, reducing propagation errors in intermediate frames. The combination of these components ensures effective generalization to unseen anatomical structures and improves temporal coherence in IUS videos. We evaluate OralSAM on both IUS and public datasets, demonstrating superior performance over state-of-the-art methods. Unlike conventional methods, our approach significantly reduces annotation effort while maintaining high segmentation accuracy. Our approach provides a scalable solution for real-time clinical applications, enabling more efficient and accurate periodontal disease assessment. Code is available at <https://github.com/BioMedCom/OralSAM>.

**Keywords:** Intraoral ultrasound videos · Periodontal disease · Segment anything model · Segmentation · One-shot learning.

## 1 Introduction

Periodontal disease is a prevalent oral condition that affects the structural integrity of the periodontium, leading to inflammation, gingival recession, and ultimately, tooth loss [11,18]. Intraoral ultrasound (IUS) has emerged as a promising imaging modality for assessing periodontal structures, offering advantages over traditional techniques such as cone beam computed tomography and radiography [1,13]. IUS is non-invasive, provides real-time imaging, is portable, operates without radiation, and is cost-effective, making it well-suited for clinical applications. Accurate segmentation of periodontal structures in IUS videos is essential for early diagnosis, disease progression monitoring, and treatment planning [16]. Despite these advantages, existing segmentation methods rely heavily on fully supervised deep learning, requiring large-scale manual annotations, which are labor-intensive and impractical for clinical deployment [17,7,12,15]. Moreover, most approaches focus on static image segmentation, failing to account for temporal coherence in ultrasound videos [6,5].

Recent advancements in foundation models, particularly the Segment Anything Model (SAM) [4], have demonstrated exceptional performance in zero-shot image segmentation. However, applying SAM directly to medical images, especially ultrasound, remains challenging due to unique imaging constraints such as low contrast, speckle noise, and anatomical variability [14]. While domain-specific adaptations of SAM, such as MedSAM [10], SAMed [19], and SAMUS [9], have attempted to bridge this gap, they still suffer from key limitations. SAMUS, for instance, is specifically designed for static ultrasound segmentation but does not incorporate temporal information necessary for video segmentation. These methods continue to require extensive manual prompts, lack robust temporal modeling, and struggle with generalizing to unseen anatomical structures. The reliance on user-provided prompts also reduces their applicability in real-time clinical workflows. A major limitation of existing medical SAM-based models is their inability to incorporate temporal information [14], which is crucial for video segmentation. Current adaptations focus primarily on static frames, leading to inconsistent segmentation across consecutive frames. This results in unreliable predictions that limit their usability in real-time medical imaging applications. Additionally, these models still demand dense annotations for effective training, restricting their scalability and deployment potential in clinical settings.

To address these challenges, we propose OralSAM, a one-shot segmentation network specifically designed for IUS videos. Unlike existing adaptations of SAM, our approach integrates temporal information through an adaptive feature correlation (AFC) module, ensuring segmentation consistency across video frames. Additionally, we introduce a self-prompting strategy that eliminates the need for expert-provided point prompts. Our method requires only a single ground truth annotation for the first frame, leveraging optical flow-based motion cues to dynamically adjust segmentation prompts, leading to more robust and stable predictions over time. Our approach is further strengthened by a self-correction mechanism, which refines mask embeddings adaptively, reducing error propagation across frames. This innovation ensures that even in

challenging low-contrast ultrasound environments, our method maintains high segmentation accuracy with minimal manual intervention. Through extensive evaluations on IUS and public datasets, we demonstrate that OralSAM outperforms state-of-the-art recent medical foundation models, achieving higher Dice and Intersection over Union (IoU) scores while significantly reducing annotation dependency. By addressing the limitations of existing medical foundation models, our study represents a significant step toward the practical deployment of AI-driven segmentation tools in clinical workflows. Our approach enables more efficient and accurate segmentation in real-time medical imaging applications, facilitating broader adoption of foundation models in healthcare and improving diagnostic efficiency and accessibility.

## 2 Method

The exceptional performance of SAM is largely attributed to its robust prompts and the strong feature representations extracted by the image encoder. To leverage this in our one-shot segmentation network, we utilize temporal cues within videos, enabling segmentation using only a single-frame annotation and a ground-truth prompt. We integrate an AFC module with a self-point prompting mechanism to enhance segmentation accuracy and ensure consistency across frames. The overall illustration of our one-shot intraoral video segmentation network is depicted in Fig. 1.

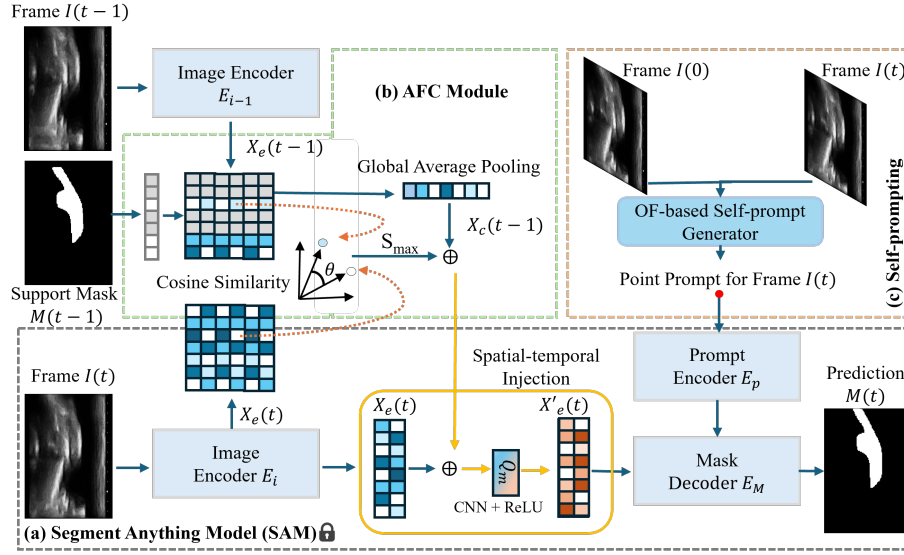


Fig. 1. Illustration of the proposed OralSAM network.

## 2.1 Adaptive Feature Correlation Module

To effectively model the temporal coherence of anatomical structures in IUS videos, we introduce the Adaptive Feature Correlation Module (AFCM). Given an input query frame  $I_t$  and support frame  $I_{t-1}$ , we extract their high-level representations using the Image Encoder  $E_i$  and  $E_{i-1}$ , respectively:

$$X_e(t) = E_i(I_t), \quad X_e(t-1) = E_{i-1}(I_{t-1}). \quad (1)$$

The support mask  $M(t-1)$  represents the predicted segmentation mask from the previous frame and is upsampled using bilinear interpolation to match the dimensions of the corresponding encoder feature map  $X_e(t-1)$ . This ensures that both tensors are spatially aligned for element-wise operations. The upsampled mask highlights the most relevant regions of the feature map by performing a Hadamard product (element-wise multiplication), resulting in a refined feature map  $X'_e(t-1) = X_e(t-1) \odot M(t-1)$ , where  $\odot$  denotes the Hadamard product.

**Temporal Feature Correlation** To measure the relevance between the query and support features, we compute a cosine similarity map:

$$S = \frac{X'_e(t-1) \cdot X_e(t)}{|X'_e(t-1)| \cdot |X_e(t)| + \epsilon}, \quad (2)$$

where  $\epsilon = 10^{-7}$  is a small constant for numerical stability. The similarity map is then refined by extracting its maximum response across feature channels,  $S_{max} = \max_c(S)$ . To further enhance feature consistency, we apply a weighted global average pooling operation using the previous frame's predicted mask  $M_{t-1}$ :

$$X_{t-1}^c = \frac{\sum_{i,j} X_e(t-1)(i,j) M_{t-1}(i,j)}{\sum_{i,j} M_{t-1}(i,j) + \delta}, \quad (3)$$

where  $\delta$  is a smoothing factor set to  $5 \times 10^{-4}$  and  $i, j$  index the spatial positions over the height and width of the feature map. The refined feature representation  $X_{t-1}^c$  is concatenated with the current frame feature  $X_e(t)$  and the similarity map  $S_{max}$  along the channel dimension. The combined tensor is then passed through  $Q_m(\cdot)$ , a lightweight convolutional fusion function implemented using a  $1 \times 1$  convolution followed by a ReLU activation. This operation generates enriched feature embeddings that enhance temporal alignment and ensure segmentation consistency across frames. This adaptive correlation strategy enables our method to achieve robust segmentation in ultrasound video sequences with minimal supervision.

## 2.2 Optical-Flow-Based Self-Prompt Generation Strategy

SAM-based segmentation frameworks rely on expert-provided prompts, but in one-shot ultrasound video segmentation, manually annotating each frame is impractical due to non-rigid anatomical transformations caused by freehand probe

movement. A single prompt in the first frame does not naturally extend across frames, as the target structure undergoes non-uniform motion, leading to segmentation drift. To address this, we propose an optical flow-based self-prompting strategy that dynamically propagates point prompts based on motion cues, ensuring consistent segmentation.

**Optical Flow Estimation** Given a ground-truth segmentation mask  $M_0$  from the first frame  $I_0$ , we extract its contour points  $P_0$ , which serve as key features for tracking anatomical structures. Instead of selecting arbitrary keypoints, we leverage these anatomical contours to define the region of interest. To estimate motion between  $I_0$  and  $I_t$  ultrasound frames, we compute sparse feature correspondences using the Lucas-Kanade pyramidal optical flow method [2,20],

$$P_t = \text{calcOpticalFlowPyrLK}(I_0, I_t, P_0). \quad (4)$$

To mitigate drift, we enforce a spatial consistency constraint: if the mean squared distance between tracked points and their original cluster center exceeds a threshold  $\delta$ , the points are reset to their initial locations. This helps maintain anatomical accuracy and prevents points from deviating from the desired structure. Using the valid correspondences  $(P_0, P_t)$ , we estimate an affine transformation matrix  $A$  that aligns the mask  $M_0$  with the current frame. This transformation matrix is then applied to warp the segmentation mask  $M_0$  onto the current frame:

$$M_t^w = \text{warpAffine}(M_0, A). \quad (5)$$

Once the segmentation mask is warped onto the next frame, we generate self-point prompts for subsequent frames without requiring additional expert input. Instead of selecting arbitrary points, we sample prompt locations directly from the warped mask:  $P_t^* = \text{RandomSample}(M_t^w)$ , where  $P_t^*$  is the selected point prompt for SAM-based segmentation in the next frame. This automatic point selection ensures that segmentation remains continuous across frames while reducing temporal drift.

### 2.3 Loss Function

**Self-Correction Mechanism** To further improve segmentation consistency, we integrate a self-correction mechanism inspired by knowledge distillation [3]. This method refines predictions by leveraging both past and present segmentation results. To implement this, we compute the probability distributions  $p_t$  and  $q_t$  across consecutive frames from the output logits  $\hat{y}_{t+1}$  and  $\hat{y}_t$  respectively:

$$p_t = \text{softmax}\left(\frac{\hat{y}_{t+1}}{T}\right), \quad q_t = \text{softmax}\left(\frac{\hat{y}_t}{T}\right) \quad (6)$$

where  $p_t$ ,  $q_t$  corresponds to the softmax of the logits at time  $t + 1$  and  $t$  respectively, and  $T$  is the temperature parameter controlling the sharpness of the

probability distribution, which encourages consistency between sequential frame predictions. The KL divergence loss function is implemented as:

$$L_{KL} = \text{KLDivLoss}(\log p_t, q_t) \times T^2 \quad (7)$$

where  $\text{KLDivLoss}$  denotes the Kullback-Leibler divergence loss computed in a batch-wise manner, ensuring robust and stable segmentation alignment over time. During training, the overall loss function integrates the segmentation loss (combining cross-entropy and Dice losses) with the self-correction terms:

$$L = (1 - \lambda)L_{CE} + \lambda L_{Dice} + \beta L_{KL} \quad (8)$$

where  $\lambda$  balances the segmentation losses, and  $\beta$  controls the influence of self-correction mechanisms. This ensures that predictions improve over time without additional computational overhead during inference.

## 2.4 Implementation Details

We built our network based on SAMUS, which adopts the standard ViT-B/16 architecture as the segmentation backbone. The model was trained on the Narval Compute Canada cluster with A100 GPUs, utilizing CUDA 12.2. We employed the AdamW optimizer with an initial learning rate of  $1 \times 10^{-4}$ , a batch size of 1. The confidence threshold  $\tau$  was set to 0.5. The loss weight coefficients  $\lambda, \beta$  were set to 0.8, and 0.002, respectively. The training was conducted for 100 epochs using grayscale ultrasound video frames resized to  $256 \times 256$ . For training our one-shot segmentation network, we utilized only the first frame’s manual annotation of each video as ground truth for loss computation.

# 3 Experiments and Results

## 3.1 Dataset and Evaluation Metrics

This study used IUS data collected from 31 orthodontic patients (25 females and 6 males) aged 11 to 57 years [6]. Ethical approval was obtained from the Research Ethics Committee of the University of Alberta, and formal consent was secured from all patients and their legal guardians. Data acquisition was conducted using a custom-built handheld 20-MHz intraoral transducer system, with the transducer positioned at the midline of each tooth to capture MP4 video recordings at an imaging depth of 13 mm. The dataset includes ultrasound scans of both mandibular and maxillary teeth, comprising 63 mandibular incisors, 43 mandibular canines, 26 mandibular first premolars, 51 maxillary incisors, 28 maxillary canines, and 27 maxillary first premolars. For one-shot segmentation evaluation, a random subset of 9 or 10 frames was selected from each video. Gingival structures were manually annotated by an imaging scientist with nine years of experience in US imaging, and all ground truth labels were subsequently verified by an experienced orthodontist. Additionally, to evaluate the generalizability of the proposed method, we assessed its performance on the publicly available CAMUS echocardiography dataset [8].

### 3.2 Comparison with the state-of-the-art foundation models

To ensure a fair comparison with recent state-of-the-art foundation models, we implemented their methods using the official repositories while maintaining the same backbone architecture. Table 1 presents the quantitative results of the


**Table 1.** Quantitative comparison on IUS and CAMUS dataset.

Method	Intraoral		CAMUS	
	Dice (%)	IoU (%)	Dice (%)	IoU (%)
SAM [4]	$0.79 \pm 0.08$	$0.67 \pm 0.11$	$0.88 \pm 0.05$	$0.80 \pm 0.08$
SAMUS [9]	$0.87 \pm 0.07$	$0.79 \pm 0.09$	$0.92 \pm 0.03$	$0.86 \pm 0.05$
SAMed [19]	$0.87 \pm 0.03$	$0.81 \pm 0.05$	$0.91 \pm 0.03$	$0.89 \pm 0.03$
MedSAM [10]	$0.89 \pm 0.04$	$0.82 \pm 0.07$	$0.92 \pm 0.02$	$0.86 \pm 0.04$
<b>OralSAM</b>	<b><math>0.91 \pm 0.04</math></b>	<b><math>0.83 \pm 0.06</math></b>	<b><math>0.94 \pm 0.02</math></b>	<b><math>0.89 \pm 0.02</math></b>

comparison on our IUS and CAMUS dataset. Our proposed one-shot segmentation network achieves the highest performance across both Dice and IoU metrics with a Dice coefficient of  $0.91 \pm 0.04$ , surpassing all other models. The second-best performance is observed with MedSAM, which attains a Dice score of  $0.89 \pm 0.04$ . Although our method offers marginal improvements in a one-shot setting, it consistently outperforms all competing approaches across both evaluation metrics. To further evaluation, we provide qualitative visualizations in Fig. 2, highlighting challenging cases. In each row, the first figure presents the manual segmentation annotation overlaid in red. In particular, SAM struggles with significant segmentation errors, especially in low-contrast regions, leading to misaligned contours. SAMUS and SAMed demonstrate inconsistent boundary delineation, often misidentifying anatomical structures. While MedSAM produces smoother and more precise contours, minor deviations from the ground truth persist. In contrast, our proposed method consistently delivers the most accurate segmentation, closely following anatomical structures. As evident from Fig. 2, the proposed method consistently delivers the most accurate segmentation, closely following anatomical structures and effectively minimizing errors in both over-segmentation and under-segmentation scenarios.

### 3.3 Ablation study

To assess the contribution of each component in the proposed network, we conducted ablation studies using the IUS dataset. Table 2 presents the Dice and IoU scores obtained from different model configurations. The baseline model achieves a Dice score of  $0.87 \pm 0.09$  and an IoU of  $0.79 \pm 0.11$ , serving as a reference for subsequent improvements. Incorporating the AFC module enhances performance, yielding a Dice score of  $0.88 \pm 0.08$  and an IoU of  $0.80 \pm 0.11$ , indicating its effectiveness in improving segmentation accuracy across video frames. Further integrating optical-flow-based self-point prompt (SPP) generation strategy



**Fig. 2.** Visual comparison with state-of-the-art (SOTA) methods on the IUS test set. Red and green regions represent the ground truth and model predictions, respectively. Yellow boxes highlight areas where the OralSAM predictions closely align with the ground truth, while red boxes indicate regions with noticeable discrepancies.

automates self-point prompting based on ultrasound frame movement, thereby reducing the need for extensive expert-provided prompts. The proposed model, which integrates all components, achieves the highest performance with a Dice score of  $0.91 \pm 0.04$  and an IoU of  $0.83 \pm 0.06$ . This highlights the effectiveness of our architectural enhancements in improving segmentation accuracy while minimizing errors in boundary delineation.

**Table 2.** Ablation study on the IUS testing dataset.

Method	Dice (%)	IoU (%)
Baseline	$0.87 \pm 0.09$	$0.79 \pm 0.11$
Baseline + AFC	$0.88 \pm 0.08$	$0.80 \pm 0.11$
Baseline + AFC + SPP	$0.90 \pm 0.04$	$0.82 \pm 0.07$
<b>OralSAM</b>	<b><math>0.91 \pm 0.04</math></b>	<b><math>0.83 \pm 0.06</math></b>



## 4 Conclusion

We propose OralSAM, a one-shot segmentation network for intraoral ultrasound imaging. By integrating temporal information and self-point prompting, our method enables automated segmentation with minimal expert intervention. Evaluations show that OralSAM outperforms existing models in segmentation accuracy and consistency. Our framework reduces reliance on manual annotations, improving feasibility for real-time clinical deployment. This scalable approach supports broader adoption of foundation models in medical imaging, improving diagnostic workflows and clinical accessibility.

**Acknowledgments.** The authors gratefully acknowledge Alberta Innovates, Canada, for the generous support of graduate studentships (L. Kumaralingam and A. Sivaanpu). This study was also supported by the National Research Council of Canada Industrial Research Assistance Program (IRAP), MITACS, and Alberta Innovates Accelerating Innovations into CarE (AICE) program (Grant No: RES0056222).

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

## References

1. Chifor, R., Badea, M.E., Mitrea, D.A., Badea, I.C., Crisan, M., Chifor, I., Avram, R.: Computer-assisted identification of the gingival sulcus and periodontal epithelial junction on high-frequency ultrasound images. *Medical ultrasonography* **17**(3), 273–279 (2015)
2. Hallock, L.A., Sud, B., Mitchell, C., Hu, E., Ahamed, F., Velu, A., Schwartz, A., Bajcsy, R.: Toward real-time muscle force inference and device control via optical-flow-tracked muscle deformation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* **29**, 2625–2634 (2021)
3. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)
4. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. In: 2023 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 3992–4003 (2023). <https://doi.org/10.1109/ICCV51070.2023.00371>
5. Kumaralingam, L., Dinh, H.B., Nguyen, K.C.T., Punithakumar, K., Kaipatur, N.R., Lou, E.H., Major, P.W., Le, L.H.: Fully-automated alveolar bone level measurements in adolescents via landmark localization in intraoral ultrasound videos. In: 2023 IEEE International Ultrasonics Symposium (IUS). pp. 1–4. IEEE (2023)
6. Kumaralingam, L., Dinh, H.B., Nguyen, K.C.T., Punithakumar, K., La, T.G., Lou, E.H., Major, P.W., Le, L.H.: Detsegdiff: A joint periodontal landmark detection and segmentation in intraoral ultrasound using edge-enhanced diffusion-based network. *Computers in Biology and Medicine* **182**, 109174 (2024)
7. Le, L.H., Nguyen, K.C.T., La, T.G., Nguyen, V.D., Le, M.B., Kumaradevan, P., Kaipatur, N., Major, P.W., Lou, E.H.: Intraoral ultrasound imaging using a rotational transducer with periodontal feature identification by machine learning. *ACS sensors* **9**(8), 3898–3906 (2024)

8. Leclerc, S., Smistad, E., Pedrosa, J., Østvik, A., Cervenansky, F., Espinosa, F., Espeland, T., Berg, E.A.R., Jodoin, P.M., Grenier, T., Lartizien, C., D'hooge, J., Lovstakken, L., Bernard, O.: Deep learning for segmentation using an open large-scale dataset in 2d echocardiography. *IEEE Transactions on Medical Imaging* **38**(9), 2198–2210 (2019). <https://doi.org/10.1109/TMI.2019.2900516>
9. Lin, X., Xiang, Y., Zhang, L., Yang, X., Yan, Z., Yu, L.: Samus: Adapting segment anything model for clinically-friendly and generalizable ultrasound image segmentation. *arXiv preprint arXiv:2309.06824* (2023)
10. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. *Nature Communications* **15**(1), 654 (2024)
11. Nazir, M., Al-Ansari, A., Al-Khalifa, K., Alhareky, M., Gaffar, B., Almas, K.: Global prevalence of periodontal disease and lack of its surveillance. *The Scientific World Journal* **2020**(1), 2146160 (2020)
12. Nguyen, K.C.T., Duong, D.Q., Almeida, F.T., Major, P.W., Kaipatur, N.R., Pham, T.T., Lou, E.H., Noga, M., Punithakumar, K., Le, L.H.: Alveolar bone segmentation in intraoral ultrasonographs with machine learning. *Journal of dental research* **99**(9), 1054–1061 (2020)
13. Nguyen, K.C.T., Le, L.H., Kaipatur, N.R., Zheng, R., Lou, E.H., Major, P.W.: High-resolution ultrasonic imaging of dento-periodontal tissues using a multi-element phased array system. *Annals of biomedical engineering* **44**, 2874–2886 (2016)
14. Painchaud, N., Duchateau, N., Bernard, O., Jodoin, P.M.: Echocardiography segmentation with enforced temporal consistency. *IEEE Transactions on Medical Imaging* **41**(10), 2867–2878 (2022)
15. Pan, Y.C., Chan, H.L., Kong, X., Hadjiiski, L.M., Kripfgans, O.D.: Multi-class deep learning segmentation and automated measurements in periodontal sonograms of a porcine model. *Dentomaxillofacial Radiology* **51**(3), 20210363 (2022)
16. Perry, D.A., Beemsterboer, P.L., Essex, G.: *Periodontology for the Dental Hygienist-E-Book*. Elsevier Health Sciences (2015)
17. Qi, B., Sasi, L., Khan, S., Luo, J., Chen, C., Rahmani, K., Jahed, Z., Jokerst, J.V.: Machine learning for automated identification of anatomical landmarks in ultrasound periodontal imaging. *Dentomaxillofacial Radiology* p. twaf001 (2025)
18. Scannapieco, F.A., Gershovich, E.: The prevention of periodontal disease—an overview. *Periodontology 2000* **84**(1), 9–13 (2020)
19. Zhang, K., Liu, D.: Customized segment anything model for medical image segmentation. *arXiv preprint arXiv:2304.13785* (2023)
20. Zhang, Q., Yang, D., Zhu, Y., Liu, Y., Ye, X.: An optimized optical-flow-based method for quantitative tracking of ultrasound-guided right diaphragm deformation. *BMC Medical Imaging* **23**(1), 108 (2023)