

Anatomically Consistent TMJ Disc Segmentation via Semantic Anchoring and Clinical Priors

Dayun Ju¹ Chanyoung Kim^{1,2} Sunyoung Jung¹ Hyo-Jung Jung³
Chena Lee³ Younjung Park³ Seong Jae Hwang^{†1}

¹ Department of Artificial Intelligence, Yonsei University, Seoul, Republic of Korea

² Department of Computer Science, Emory University, Atlanta, Georgia, USA

³ College of Dentistry, Yonsei University, Seoul, Republic of Korea

{juda0707, chanyoung, sunyoungj, seongjae}@yonsei.ac.kr,
chanyoung.kim@emory.edu,

{hjjung, chenalee, darkstar}@yuhs.ac

Abstract. Segmenting the temporomandibular joint (TMJ) disc from MRI is essential for accurate diagnosis of internal derangement, yet it remains unreliable in practice due to its small size, low contrast, and morphological variability. Existing methods, primarily adapted from general segmentation architectures, often produce fragmented or anatomically inconsistent masks, leading to unstable measurements of disc position and shape for downstream diagnosis. To address these challenges, we propose TISC, a TMJ disc segmentation framework that integrates semantic anchoring with clinical metadata-guided boundary refinement. The framework first establishes robust disc localization in the foundation model feature space via a Prototypical Semantic Anchoring (PSA) module that aggregates adjacent-slice MedDINOv3 features and derives a prototype-driven similarity map. It then performs targeted boundary refinement through a Clinical-Metadata Point Refinement (C-MPR) module, with point-wise predictions modulated by Mouth Open Limitation (MOL), a clinical indicator associated with disc displacement without reduction. On a large-scale cohort of 2,488 PD MRI volumes from 1,300 patients, our method achieves up to a 4.96 Dice improvement over strong baselines across diverse architectures, delivering more anatomically coherent and clinically reliable TMJ disc segmentation.

Keywords: TMJ disc segmentation · Prototype-guided localization · Clinical metadata

1 Introduction

Temporomandibular joint (TMJ) disc is a vital fibrocartilaginous structure that ensures smooth mandibular movement and optimal load distribution [18]. Accurate interpretation of the disc is essential for managing internal derangement, which affects up to 12% of the population and is a leading cause of chronic

[†] Corresponding Author

orofacial pain [23]. Although Magnetic Resonance Imaging (MRI) is the gold standard for diagnosis, reliable assessment of disc displacement and structural abnormalities remains challenging, often requiring highly specialized radiological expertise to resolve ambiguities between the disc and surrounding soft tissues [2, 13]. If automated segmentation were available, it could facilitate objective and reproducible quantification of disc displacement and morphology, enabling more standardized treatment planning. However, this task is particularly difficult, as the disc typically occupies only about 0.2% of the total pixel area on an MRI slice and exhibits low contrast relative to adjacent ligaments and effusions [17].

In the TMJ domain, various data-driven approaches have been explored for diagnostic modeling and clinical assessment [9, 10, 14]. Among these efforts, TMJ disc segmentation has primarily been studied through task-specific adaptations of general segmentation architectures, such as UNet-based models [4, 5, 8, 21]. More recently, medical foundation models, including MedSAM [16] have also been applied to medical image segmentation. Despite their broad utility, general-purpose segmentation models struggle in TMJ disc segmentation for two primary reasons. First, reliable localization is inherently difficult because the disc is extremely small and exhibits low contrast, providing limited semantic cues [17]. This problem becomes more pronounced in pathological conditions such as disc displacement without reduction, where the disc deviates from its typical anatomical position. Consequently, spatial uncertainty increases, weakening positional cues and making models prone to confusion with adjacent anatomy or imaging artifacts, leading to false positives. Second, even when roughly localized, precise boundary delineation remains challenging. The disc lacks a consistent shape and may appear folded, compressed, or morphologically distorted across disease states, while subtle intensity transitions with neighboring tissues further obscure its margins. These factors often result in ambiguous boundaries and over-segmentation into surrounding structures [17].

Motivated by these challenges, we propose TISC, a TMJ disc segmentation framework that integrates semantic anchoring with metadata-guided boundary refinement. The framework first establishes disc localization in the foundation model feature space. To this end, the Prototypical Semantic Anchoring (PSA) module aggregates adjacent-slice MedDINOv3 [15] features and employs a cross-attention query extractor to generate a similarity-based localization map at the bottleneck, stabilizing predictions and reducing false positives under pathological displacement. Building upon this disc-focused representation, the Clinical-Metadata Point Refinement (C-MPR) module adaptively adjusts boundary predictions at the point level [11] conditioned on clinical metadata, enabling correction of ambiguous regions. We further incorporate Mouth Open Limitation (MOL), a binary indicator of limited mouth opening (<40 mm) [23], as auxiliary metadata due to its strong association with disc displacement without reduction (DDwoR), a condition characterized by abnormal disc position and deformation [7, 12, 20, 22]. Treated as an indicator of displacement severity, MOL guides spatial attention in PSA and modulates refinement features via FiLM layers, promoting adaptive localization and boundary correction across pathological states.




Fig. 1: Overview of TISC. (a) The PSA module establishes disc localization at the bottleneck, followed by (b) the C-MPR module for boundary refinement during upscaling, enabling anatomically consistent and precise segmentation.

The proposed framework is evaluated on a large-scale in-house cohort comprising 2,488 PD MRI volumes and 4,601 manually annotated slices from 1,300 patients. Through extensive patient-level cross-validation, our model consistently outperforms state-of-the-art baselines across diverse architectures. It achieves notable improvements in both region-level and boundary metrics, producing sharper and anatomically coherent segmentation.

2 Method

As illustrated in Fig. 1, we propose TISC for TMJ disc segmentation via semantic anchoring and clinical metadata-guided mask refinement. The framework first employs a semantic anchoring module (PSA, Sec. 2.1) that distills foundation model features to establish anatomical anchors for localization. It then incorporates a clinical metadata-guided point refinement module (C-MPR, Sec. 2.2) for fine-grained boundary refinement to ensure anatomical fidelity of the mask.

Notably, we further augment this framework by integrating Mouth Open Limitation (MOL), defined as mouth opening less than 40 mm, metadata as a clinical surrogate for disc displacement without reduction (DDwoR). Clinically, a positive MOL status (MOL=1) is strongly associated with DDwoR [22, 23], a condition characterized by abnormal disc position [7, 20] and exhibiting deformed morphologies such as folded or globular shapes [12]. We leverage these insights by using MOL to adaptively scale attention temperature in the PSA module for wider spatial anchoring and to guide feature modulation via FiLM layers in the C-MPR module. This design enables the model to capture structural details and maintain anatomical consistency across diverse disease states.




Fig. 2: Visualization of similarity-based disc localization maps. From left to right, the figure shows the original MRI slice, the ground truth disc mask, the MedDINOv3 feature map, the similarity map using M_{GT} , the similarity map from the learned prototype, and the final localization map after temperature scaling.

2.1 Disc-localized Multi-context Encoding

To exploit high-fidelity foundation model features for anatomical structures, we propose the Prototypical Semantic Anchoring (PSA) module. PSA reinforces disc signals via (I) volumetric context and (II) distills anatomical identity through a learnable query in the MedDINOv3 [15] space. (III) Projecting this identity back to the spatial domain establishes a precise anchor to guide segmentation.

(I) Multi-context Feature Aggregation. To incorporate volumetric context, we aggregate features from a target slice s_i and its adjacent slices $s_{i\pm 1}$. Each slice s_j ($j \in \{i-1, i, i+1\}$) is encoded by a pre-trained MedDINOv3 encoder to obtain feature maps $F_j \in \mathbb{R}^{L \times C}$, where L denotes the number of spatial tokens and C the embedding dimension. We then compute an aggregated representation as $\hat{F} = (F_{i-1} + 2F_i + F_{i+1})/4$. This aggregation enforces inter-slice consistency and stabilizes disc representations against slice-level ambiguity and noise.

(II) Semantic Prototype Extraction. Using the aggregated feature map $\hat{F}$, the PSA module distills the disc-representative semantic features into a latent prototype Φ to determine *what* constitutes the target, providing a semantic foundation for modulating the UNet bottleneck. To achieve this, we define a single learnable query $Q \in \mathbb{R}^C$ that selectively aggregates spatial information via a cross-attention mechanism: $\Phi = \text{Softmax}\left(\frac{Q\hat{F}^T}{\sqrt{C}}\right)\hat{F}$, where $\hat{F}$ acts as a key and value. This operation selectively pools features across the spatial dimension L to form a latent prototype that contains the disc’s semantic properties. To spatially align the prototype Φ , we optimize PSA with a spatial alignment loss $\mathcal{L}_{anchor} = \text{BCE}(\text{Norm}(S), M_{GT})$, where BCE denotes the Binary Cross Entropy loss. Here, S is the cosine similarity map between Φ and the aggregated feature map $\hat{F}$, and $\text{Norm}(\cdot)$ normalizes S to $[0, 1]$. The ground truth mask M_{GT}

is Gaussian-smoothed to prevent overly sharp supervision and encourage spatial tolerance. This ensures the query-based filter remains focused on the disc’s identity even under complex imaging artifacts.

(III) Disc Localization and Bottleneck Modulation. Given the similarity map S , we project the distilled semantic identity onto the spatial feature grid to determine where feature modulation should occur. We obtain a Disc Localization Map A by applying a temperature-scaled spatial softmax to S as

$$A(y, x) = \frac{\exp(S(y, x)/\tau)}{\sum_{y', x'} \exp(S(y', x')/\tau)}, \quad (1)$$

where (y, x) are spatial indices. The temperature τ is modulated by MOL, as illustrated in Fig. 2. Specifically, $\tau = 0.05$ for MOL=0 yields a sharper attention map over stable anatomical regions, whereas $\tau = 0.1$ produces a broader distribution to accommodate pathological displacement, reflecting the higher likelihood of disc displacement when MOL=1. The resulting map A acts as a semantic spatial mask that modulates $\hat{F}$ via element-wise multiplication, suppressing non-disc responses. The localized features are then injected into the UNet bottleneck through an additive skip connection to guide decoding toward disc-relevant regions.

2.2 Metadata-Augmented Point Refinement

To enable fine-grained adjustments at the point level [11], we introduce a Clinical-Metadata-Guided Point Refinement (C-MPR) module. By leveraging MOL as a clinical prior, C-MPR reduces morphological ambiguities and enables high-fidelity boundary reconstruction, particularly in cases of severe pathological displacement. This module operates through two main stages: (I) Progressive Point-based Upsampling and (II) Metadata-Driven Feature Modulation.

(I) Progressive Point-based Upsampling. The C-MPR module performs boundary refinement starting from a coarse logit map $M_{coarse} \in \mathbb{R}^{1 \times H/2 \times W/2}$ predicted by the decoder. To preserve global structural coherence, M_{coarse} is supervised with an auxiliary loss $\mathcal{L}_{aux}$ combining BCE and Soft Dice. Unlike existing uncertainty-driven refinement [11], C-MPR adopts a semantic-guided point sampling strategy. We sample P candidate points by combining (i) uncertainty-based selection near the decision boundary, where coarse probabilities approach 0.5, and (ii) prototype-aware sampling from regions with high similarity responses in S . Approximately 25% of points are drawn from semantically salient regions, promoting refinement around disc-relevant structures. For each selected point, a feature vector h is constructed by concatenating coarse logits with high-resolution encoder features $E_{fine} \in \mathbb{R}^{C_e \times H \times W}$. This design enables localized boundary correction guided by prediction uncertainty and semantic identity.

Table 1: Quantitative Analysis. Results are reported as mean(standard deviation) over 5-fold cross-validation. Performance is evaluated using Dice, clDice, Detection rate, HD95, and ASSD. [†] indicates $p < 0.05$ compared to baselines.

Method	Region-level (%) $\uparrow$			Boundary (mm) $\downarrow$	
	Dice	clDice	Detection	HD95	ASSD
UNETR [5]	58.96(0.38)	63.58(0.71)	81.66(1.59)	13.51(0.40)	2.96(0.07)
MedSAM [16]	53.30(0.51)	63.11(0.94)	69.73(2.01)	10.33(0.07)	2.68(0.03)
UNet [†] [21]	76.58(0.33)	86.23(0.70)	97.41(1.22)	7.19(0.38)	1.37(0.07)
+ Ours	80.46(0.36)	89.94(0.58)	98.85(0.38)	5.23(0.16)	1.02(0.04)
nnUNet [†] [8]	78.02(0.22)	87.54(0.33)	98.41(0.53)	6.33(0.34)	1.23(0.03)
+ Ours	80.12(0.96)	89.77(1.14)	98.62(0.45)	5.17(0.35)	1.27(0.28)
SwinUNETR [†] [4]	75.34(0.41)	84.21(0.62)	96.52(0.82)	7.69(0.25)	1.57(0.05)
+ Ours	80.30(0.35)	89.67(0.54)	99.00(0.27)	5.16(0.24)	1.03(0.07)

(II) Metadata-Driven Feature Modulation. To handle morphology changes under pathological conditions, we process point-wise features h using a FiLM-based MLP [19] conditioned on MOL. While PSA provides a global anatomical anchor at the bottleneck, h captures local multi-scale context, enabling C-MPR to adapt boundary refinement based on clinical metadata. Each FiLM layer applies a channel-wise affine transformation

$$h' = \gamma(\text{MOL}) \odot h + \beta(\text{MOL}), \quad (2)$$

where $\gamma(\text{MOL})$ and $\beta(\text{MOL})$ denote scaling and shifting vectors generated from a learnable embedding of MOL. This conditioning adaptively reweights feature responses, increasing sensitivity to structural deformations in MOL=1 cases while maintaining stable boundary delineation in MOL=0 cases. The refined predictions are supervised with a point-wise BCE loss $\mathcal{L}_{point}$.

Total Loss. The framework is optimized via a multi-task objective function:

$$\mathcal{L}_{total} = \mathcal{L}_{seg} + \lambda_{point}\mathcal{L}_{point} + \lambda_{aux}\mathcal{L}_{aux}. \quad (3)$$

where $\mathcal{L}_{seg}$ combines BCE and soft dice loss. The weighting factors λ_{point} and λ_{aux} regulate the contributions of $\mathcal{L}_{point}$ and $\mathcal{L}_{aux}$, respectively.

3 Experiments

3.1 Experimental Settings

Dataset and Evaluation Metrics. The proposed framework is evaluated on a private dataset of 2,488 PD-weighted MRI volumes from 1,300 patients, following Institutional Review Board approval and the Declaration of Helsinki [1].




Fig. 3: Qualitative Results. Segmentation results of our model are compared against MedSAM [16], UNet [21], nnUNet [8], and SwinUNETR [4]. Yellow regions represent predicted masks, and the red boundary denotes the ground truth.

Images are acquired with a 3.0 T GE Pioneer scanner using specific parameters: TE/TR of 50/2084 ms, a 130×130 mm FOV, and a 2.5 mm slice thickness. For ground truth segmentation masks, 4,601 sagittal slices are annotated and verified through specialist consensus. The dataset also includes patient-level MOL and expert-annotated disc displacement labels. We conduct patient-level 5-fold cross-validation to ensure robustness and prevent data leakage. Performance was evaluated using Dice [3], cDice [24], and Detection (IoU > 0.3) for region-level performance, along with 95th percentile Hausdorff Distance (HD95) and Average Symmetric Surface Distance (ASSD) to measure boundary accuracy.

Implementation Details. To ensure reproducibility and consistency, a uniform hyperparameter configuration was maintained across all cross-validation folds. The weighting coefficients λ_{point} and λ_{aux} are set to 20 and 0.8, respectively. The PSA module was pre-trained for 20 epochs and subsequently frozen during segmentation training. For the full pipeline, a warm-up was performed for 10 epochs to stabilize the encoder before activating the C-MPR module. All images and corresponding labels are used at their original resolution of 512×512 .

3.2 Evaluation Results

Quantitative Evaluation. Table 1 presents the average 5-fold cross-validation results. Our framework consistently improves both region-level and boundary metrics over state-of-the-art baselines. When integrated with UNet [21], the proposed method achieves the highest Dice score of 80.46%, outperforming the original UNet (76.58%). Substantial improvements are also observed in boundary metrics, with HD95 reduced by 32.9% on the SwinUNETR baseline [4] (7.69 to 5.16). These results highlight its robustness across diverse segmentation architectures, consistently improving both localization and boundary precision.

Table 2: Ablation study on PSA, C-MPR components and MOL metadata.

Exp.#	PSA		C-MPR	MOL	Dice $\uparrow$	HD95 $\downarrow$
	Prototype	Vol. Neighbor				
1					76.54	7.49
2	✓			✓	78.99	6.22
3	✓	✓		✓	79.55	6.01
4			✓	✓	78.75	5.30
5	✓	✓	✓		78.65	6.35
6	✓	✓	✓	✓	80.10	5.27

Qualitative Analysis. Fig. 3 presents representative qualitative comparisons between our framework (UNet-based) and baseline models. The results show that baseline methods frequently produce partial segmentations that capture only the anterior portion of the disc or exhibit oversegmentation beyond the ground truth boundary (red contour). In contrast, our framework more accurately delineates the complete disc morphology, closely adhering to the annotated boundary and preserving structural continuity across the thin intermediate zone.

3.3 Ablation Study

Table 2 shows ablation study results. Given the consistent trends across all validation folds (Sec. 3.2), we conduct ablation experiments on Fold 0 using a standard UNet [21] backbone to isolate the contribution of each module.

Effect of PSA. Exp. 2 shows that incorporating PSA with single-slice features (F) improves Dice from 76.54 to 78.99. Extending PSA to multi-slice contextual features ($\hat{F}$) in Exp. 3 further increases Dice to 79.55. These results indicate that multi-context PSA enhances disc localization.

Effect of C-MPR. Comparing Exp. 1 and Exp. 4 highlights the contribution of the C-MPR module. Incorporating C-MPR reduces HD95 from 7.49 to 5.30 (29.2% reduction) while improving Dice from 76.54 to 78.75, demonstrating enhanced boundary accuracy. These results indicate that metadata-guided point refinement improves structural continuity.

Effect of Clinical Metadata Conditioning. Removing MOL-guided conditioning from PSA and C-MPR (Exp. 5) results in a performance drop compared to the full model (Exp. 6). This indicates that incorporating MOL provides meaningful clinical cues that enhance robustness under pathological variability.

3.4 Application

Segmentation-Assisted Classification. To assess whether segmentation masks improve downstream diagnosis, we performed binary classification of normal versus disc displacement. Close-open MRI slice pairs from the same joint side are used as a 2-channel input to a ResNet18 [6] baseline. Adding predicted masks

expanded it to a 4-channel input. The vanilla model achieved an AUROC of 0.61, which increased to 0.73 with our masks, approaching the performance obtained using ground truth masks (0.74). These findings indicate that accurate disc segmentation provides meaningful structural cues for downstream diagnosis.

4 Conclusion

We present TISC, a framework for anatomically consistent TMJ disc segmentation in PD-weighted MRI that unifies semantic anchoring with clinically guided refinement. The PSA module establishes a stable localization prior in the foundation model feature space via multi-slice aggregation and prototype-based modulation, mitigating mislocalization under pathological displacement, while the C-MPR module leverages MOL metadata to preserve structural continuity in morphologically distorted cases. Extensive patient-level validation confirms consistent improvements over strong baselines. By tightly coupling foundation model semantics with clinically meaningful priors, the framework enables reliable segmentation of small, low-contrast, pathology-sensitive structures.

Acknowledgments. This work was supported in part by the IITP RS-2024-00457882 (AI Research Hub Project), IITP 2020-II201361, NRF RS-2024-00345806, NRF RS-2023-002620, and RQT-25-120390.

Disclosure of Interests. The authors have no competing interests.

References

1. Association, W.M., et al.: World medical association declaration of helsinki: ethical principles for medical research involving human subjects. *Jama* **310**(20), 2191–2194 (2013)
2. Bag, A.K., Gaddikeri, S., Singhal, A., Hardin, S., Tran, B.D., Medina, J.A., Curé, J.K.: Imaging of the temporomandibular joint: An update. *World journal of radiology* **6**(8), 567 (2014)
3. Dice, L.R.: Measures of the amount of ecologic association between species. *Ecology* **26**(3), 297–302 (1945)
4. Hatamizadeh, A., Nath, V., Tang, Y., Yang, D., Roth, H.R., Xu, D.: Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. In: *International MICCAI brainlesion workshop*. pp. 272–284. Springer (2021)
5. Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H.R., Xu, D.: Unetr: Transformers for 3d medical image segmentation. In: *WACV*. pp. 574–584 (2022)
6. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
7. Hegab, A.F., Al Hameed, H.I., Karam, K.S.: Classification of temporomandibular joint internal derangement based on magnetic resonance imaging and clinical findings of 435 patients contributing to a nonsurgical treatment protocol. *Scientific Reports* **11**(1), 20917 (2021)

8. Isensee, F., Petersen, J., Klein, A., Zimmerer, D., Jaeger, P.F., Kohl, S., Wasserthal, J., Koehler, G., Norajitra, T., Wirkert, S., et al.: nnu-net: Self-adapting framework for u-net-based medical image segmentation. *arXiv preprint arXiv:1809.10486* (2018)
9. Ito, S., Mine, Y., Yoshimi, Y., Takeda, S., Tanaka, A., Onishi, A., Peng, T.Y., Nakamoto, T., Nagasaki, T., Kakimoto, N., et al.: Automated segmentation of articular disc of the temporomandibular joint on magnetic resonance images using deep learning. *Scientific Reports* **12**(1), 221 (2022)
10. Jung, H.J., Ju, D., Kim, C., Hwang, S.J., Lee, C., Park, Y.: Multimodal deep learning with anatomically constrained attention for screening mri-detectable tmj abnormalities from panoramic images. *npj Digital Medicine* (2026)
11. Kirillov, A., Wu, Y., He, K., Girshick, R.: Pointrend: Image segmentation as rendering. In: *CVPR*. pp. 9799–9808 (2020)
12. Kurita, H., Ohtsuka, A., Kobayashi, H., Kurashina, K.: The relationship between the degree of disk displacement and ability to perform disk reduction. *Oral Surgery, Oral Medicine, Oral Pathology, Oral Radiology, and Endodontology* **90**(1), 16–20 (2000)
13. Larheim, T., Abrahamsson, A., Kristensen, M., Arvidsson, L.: Temporomandibular joint diagnostics using cbct. *Dentomaxillofacial Radiology* **44**(1), 20140235 (2015)
14. Li, M., Punithakumar, K., Major, P.W., Le, L.H., Nguyen, K.C.T., Pacheco-Pereira, C., Kaipatur, N.R., Nebbe, B., Jaremko, J.L., Almeida, F.T.: Temporomandibular joint segmentation in mri images using deep learning. *Journal of Dentistry* **127**, 104345 (2022)
15. Li, Y., Wu, Y., Lai, Y., Hu, M., Yang, X.: Meddinov3: How to adapt vision foundation models for medical image segmentation? *arXiv preprint arXiv:2509.02379* (2025)
16. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. *Nature Communications* **15**, 654 (2024)
17. Nozawa, M., Ito, H., Ariji, Y., Fukuda, M., Igarashi, C., Nishiyama, M., Ogi, N., Katsumata, A., Kobayashi, K., Ariji, E.: Automatic segmentation of the temporomandibular joint disc on magnetic resonance images using a deep learning technique. *Dentomaxillofacial Radiology* **51**(1), 20210185 (2022)
18. Okeson, J.P.: Management of Temporomandibular Disorders and Occlusion-E-Book: Management of Temporomandibular Disorders and Occlusion-E-Book. Elsevier Health Sciences (2019)
19. Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: Film: Visual reasoning with a general conditioning layer. In: *AAAI*. vol. 32 (2018)
20. Poluha, R.L., Canales, G.D.I.T., Costa, Y.M., Grossmann, E., Bonjardim, L.R., Conti, P.C.R.: Temporomandibular joint disc displacement with reduction: a review of mechanisms and clinical presentation. *Journal of applied oral science* **27**, e20180433 (2019)
21. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *MICCAI*. pp. 234–241. Springer (2015)
22. Sato, S., Sakamoto, M., Kawamura, H., Motegi, K.: Long-term changes in clinical signs and symptoms and disc position and morphology in patients with nonreducing disc displacement in the temporomandibular joint. *Journal of oral and maxillofacial surgery* **57**(1), 23–29 (1999)
23. Schiffman, E., Ohrbach, R., Truelove, E., Look, J., Anderson, G., Goulet, J.P., List, T., Svensson, P., Gonzalez, Y., Lobbezoo, F., et al.: Diagnostic criteria for

- temporomandibular disorders (dc/tmd) for clinical and research applications: recommendations of the international rdc/tmd consortium network and orofacial pain special interest group. *Journal of oral & facial pain and headache* **28**(1), 6 (2014)
24. Shit, S., Paetzold, J.C., Sekuboyina, A., Ezhov, I., Unger, A., Zhylka, A., Pluim, J.P., Bauer, U., Menze, B.H.: cldice-a novel topology-preserving loss function for tubular structure segmentation. In: *CVPR*. pp. 16560–16569 (2021)