Landmark-Guided Coarse-to-Fine Registration of Intraoral Scans and Cone-Beam CT

Alex Veronese, Matteo Lugli, Omar Carpentiero,
Kevin Marchesini, Luca Lumetti, and Federico Bolelli ✉

University of Modena and Reggio Emilia, Modena, Italy
{name.surname}@unimore.it

Abstract. Accurate registration of Intraoral Scans (IOS) with Cone-Beam Computed Tomography (CBCT) enables integration of dental crown surfaces with the tooth roots and the surrounding alveolar bone. However, IOS-CBCT registration remains challenging because of limited anatomical overlap, cross-modal differences, and unreliable correspondences. Task 2 of the MICCAI STS 2026 Challenge formulates this setting as a semi-supervised rigid registration problem, in which IOS meshes must be aligned with their corresponding CBCT volumes. For this task, we propose a landmark-driven coarse-to-fine framework in which modality-specific networks predict corresponding dental landmarks from the IOS mesh and CBCT volume. A confidence-weighted RANSAC-Kabsch procedure estimates a robust initial transformation, later refined using point-to-plane Iterative Closest Point. On the challenge validation set, the method achieved a *Mean Translation Error* of 4.988 mm and a *Mean Rotation Error* of 1.421°, ranking first on the registration leaderboard at the time of evaluation. The code is available on GitHub: https://github.com/AImageLab-zip/L2L-Registration

Keywords: IOS-CBCT registration · Dental imaging · Landmark detection · RANSAC · Iterative Closest Point

1 Introduction

Cone-Beam Computed Tomography (CBCT) and Intraoral Scans (IOS) provide complementary representations of dentoalveolar anatomy [3,6,18]. CBCT captures tooth roots and alveolar bone, but its representation of dental crowns may be degraded by limited resolution, noise, partial-volume effects, and metal artifacts. In contrast, IOS provides high-resolution surface models of the visible crowns and gingiva [4], but does not capture roots or underlying bone and may exhibit stitching errors in full-arch scans [17,23]. Accurate IOS-CBCT registration therefore enables the construction of integrated crown-root-bone models for orthodontic diagnosis, implant planning, orthognathic surgery, and longitudinal assessment of tooth movement. This task is challenging because the modalities

✉ Corresponding author: federico.bolelli@unimore.it

differ in data representation and share only a limited anatomical region: the exposed dental elements. Reliable correspondence estimation may be hindered by missing teeth, similar or symmetric tooth morphology, restorations, orthodontic appliances, metal artifacts, and incomplete scans. Prior studies have also shown that registration performance depends strongly on the accuracy of tooth segmentation and identification. Moreover, large initial pose differences may cause local surface-registration methods to converge to anatomically incorrect alignments [10,23]. Learning-based approaches face the additional limitation of scarce paired data with accurate transformations.

Conventional workflows typically estimate an initial transformation from manually selected landmarks or dental surfaces and refine the alignment using Iterative Closest Point (ICP) [2]. Although computationally efficient, ICP is sensitive to initialization, partial overlap, outliers, and repetitive dental geometry. Recent automatic methods have therefore adopted coarse-to-fine pipelines combining semantic preprocessing, robust global alignment, and local surface refinement. Kim et al. [14] combined deep learning-based tooth segmentation with RANSAC-based coarse registration [8] and ICP refinement. Jang et al. [10] incorporated individual-tooth segmentation and identification into a global-to-local framework while addressing cumulative deformation in full-arch IOS scans. More recently, point-cloud-based approaches have been investigated, including FPFH-based geometric matching [15], PointNet++ feature encoding with differentiable SVD-based rigid-transformation estimation [11], and PointNetLK-based iterative alignment [9,19]. However, these approaches may remain sensitive to cross-modal geometric discrepancies, limited anatomical overlap, and large initial pose differences.

Task 2 of the MICCAI STS 2026 Challenge formulates IOS-CBCT alignment as a semi-supervised rigid-registration problem and provides labeled and unlabeled dental data for standardized evaluation. Building on recent advances in modality-specific anatomical landmark prediction [7], we propose a landmark-guided coarse-to-fine registration framework. Our main contribution is the robust use of predicted anatomical correspondences: a confidence-weighted adaptive RANSAC-Kabsch procedure identifies a reliable landmark consensus for global alignment, and the resulting inliers guide the subsequent local surface refinement. This design reduces the influence of inaccurate landmark predictions and unrelated CBCT anatomy while improving robustness to limited overlap and large initial pose differences.

2 Method

We propose a landmark-driven pipeline for rigid IOS-CBCT registration. As illustrated in Fig. 1, the IOS and CBCT data are processed in parallel to predict anatomically corresponding dental landmarks. In the IOS branch, the input mesh is centered and oriented by the *IOS-Normalizer* module according to the reference frame required by the subsequent *IOS-Landmarks* module, which estimates the three-dimensional coordinates of the selected dental landmarks. In

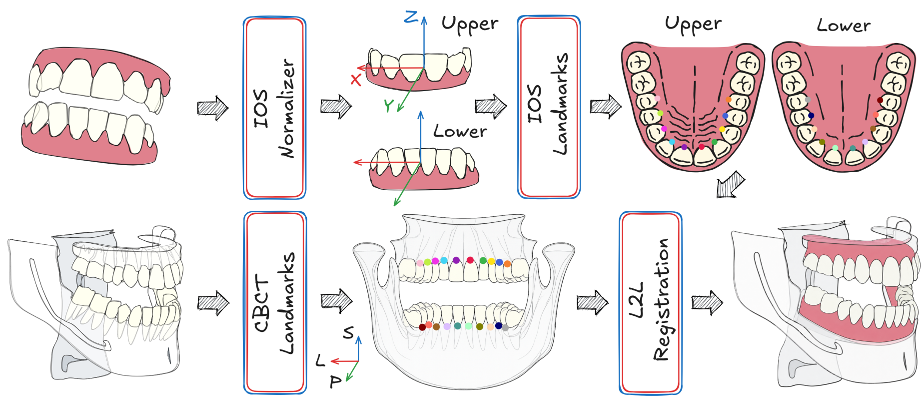

Fig. 1: Overview of the proposed IOS-CBCT registration pipeline.

parallel, the *CBCT-Landmarks* module localizes the corresponding landmarks directly in the CBCT volume. The predicted correspondences are provided to the Landmarks-to-Landmarks registration module (*L2L-Registration* in short), which estimates an initial rigid IOS-CBCT transformation using a confidence-weighted RANSAC-Kabsch procedure [8,12,13]. The resulting alignment is subsequently refined using point-to-plane ICP [2]. This coarse-to-fine strategy reduces sensitivity to large initial pose differences and improves the final alignment between the IOS surface and the corresponding CBCT dental surface.

IOS-Normalizer. Given an input dental mesh, this module[1] aligns each IOS with the reference frame used to train the IOS-Landmarks module. The normalization procedure combines principal component analysis (PCA) with a learned rotation correction. A fixed number of points is sampled from the unit-sphere normalized vertices using farthest-point sampling (FPS). PCA is then applied to the covariance matrix of the sampled point cloud. The ordered eigenvectors form an orthonormal basis B that aligns the scan with its principal geometric axes. If $\det(B) < 0$, the sign of the final eigenvector is reversed to ensure a right-handed coordinate system. Because PCA does not uniquely determine the signs of the eigenvectors, residual orientation ambiguities may remain. To resolve these ambiguities, the PCA-aligned point cloud is processed by a small PointNet-style network [5], which predicts a single corrective rotation R selected from four predefined classes: the identity rotation and 180° rotations about the x-, y-, and z-axes. The transformed mesh vertices are computed as $X' = (X - c)BR$, where c denotes the mesh centroid.

IOS-Landmarks. Given a single IOS, this module predicts dental landmarks using a two-stage deep learning architecture.[2] First, a semantic segmentation network classifies each mesh vertex into one of 17 classes, representing the 16 teeth of the input arch and a combined gingiva/noise class. At inference time,

[1] https://github.com/AImageLab-zip/IOS-Normalizer/commit/ecebe11
[2] https://github.com/AImageLab-zip/IOS-Landmarks/commit/3ee1ee6

Table 1: L2L-Registration Parameters

Stage	Parameter	Setting
(i) Landmark	Min. landmark confidence	0.25
	RANSAC hypotheses	2,000
	Landmark pairs per hypothesis	3 non-collinear pairs
	Initial inlier distance	3.5 mm
	Max. inlier distance	5.5 mm
	Initial min. consensus	4 landmarks[a]
	Fallback min. consensus	3 non-collinear landmarks
(ii) Surface	Crop margin	max(15 mm, $0.25E$)
	Marching-cubes threshold	1,500
	Max. sampled points	12,000 per surface
	Voxel size	1.0 mm
	Normal-estimation radius	3.0 mm
(iii) ICP	Max. correspondence distance	1.0 mm
	Max. iterations	40
	Relative fitness tolerance	10^{-5}
	Relative RMSE tolerance	10^{-5}

E denotes the landmark extent along the corresponding spatial axis.
[a] Applied when at least four landmark correspondences were available.
RMSE: root-mean-square error.

the tooth classes are mapped to their corresponding FDI numbers using the known arch identity (lower or upper). The vertices assigned to each tooth are processed independently by a landmark-prediction network that generates a per-vertex probability heatmap. Final landmark coordinates are obtained through thresholding and K-means clustering. Both segmentation and landmark prediction models are based on Point Transformer V3 [22]. Among the predicted landmarks, only the *Inner* ones, illustrated in Fig. 1, were used directly in the subsequent stages. This choice was motivated by the variable quality of the IOS scans observed in the training and validation sets. Some scans contained residual orthodontic brackets or dental impression material in the anterior region, potentially affecting landmarks located near the external tooth surface. To assess landmark availability across the dataset, we measured the proportion of tooth instances with a valid *Inner* prediction in each of the four hemi-arches. Applying a minimum prediction coverage threshold of 94% in each hemi-arch retained tooth positions 1–5 for subsequent processing.

CBCT-Landmarks. This module was developed using nnLandmark [7], an extension of the nnU-Net infrastructure to 3D landmark localization via heatmap regression. Landmark annotations are encoded as a $3\times3\times3$ cube around the target voxel. During inference, the input volume is processed using sliding-window prediction. For each landmark channel, the location of the maximum response defines the predicted coordinate, while the corresponding response value is retained as a confidence score. For each tooth previously selected for the IOS, only the corresponding *Inner* landmark was retained.

L2L-Registration. The final IOS-CBCT registration was performed in two stages: *(i)* a landmark-based rigid alignment, followed by *(ii)* surface-based refinement. Hyperparameters for these stages are summarized in Tab. 1. We se-

Table 2: Development environments and requirements.

Component	Configuration
Operating system	Ubuntu 22.04.5 LTS
Processor	Intel Xeon Silver 4216 CPU @ 2.10 GHz
System memory	128 GB
Graphics processor	1 × NVIDIA RTX A5000, 24 GB
CUDA version	`12.6`
Programming language	`Python 3.10`
Deep learning framework	`PyTorch 2.8`
Computer vision library	`torchvision 0.23.0`

lected the final configuration through a grid search that considered registration failures, combined translation and rotation errors, and computational cost. Predicted CBCT landmarks below the confidence threshold were excluded, and the remaining landmarks were converted from voxel indices to physical coordinates using the image metadata. An initial rigid transformation was estimated using a confidence-weighted RANSAC-Kabsch procedure [8,12,13]. For each hypothesis, three non-collinear landmark pairs were sampled and used to estimate a rigid transformation. The transformation was evaluated on all available correspondences, and the best solution was selected according to the inlier consensus, combined confidence, and residual error. When no valid consensus was obtained, the inlier threshold was progressively relaxed and, as a final fallback, the minimum consensus size was reduced to three non-collinear landmarks. The selected transformation was then re-estimated using all inliers until the consensus set stabilized and was applied to the normalized IOS mesh. For surface refinement, the landmark inliers were used to define an axis-aligned crop of the CBCT volume. A CBCT surface was extracted from the cropped region using marching cubes, after which the IOS and CBCT surfaces were sampled, voxel-downsampled, and assigned surface normals. The landmark-aligned IOS point cloud was then refined against the CBCT point cloud using point-to-plane ICP. The resulting ICP transformation was composed with the landmark-based transformation to obtain the final IOS-CBCT registration.

3 Experiments

Dataset and Evaluation Metrics. All experiments were conducted using the official MICCAI STS 2026 Challenge dataset, which comprises: *(i)* 30 cases with ground-truth 4×4 transformations and 300 unlabeled cases (*training set*) and *(ii)* 50 cases with ground-truth transformations withheld from participants and accessible only to the challenge organizers (*validation set*). Performance was evaluated using two metrics: *(i) Mean Translation Error* (MTE) and *(ii) Mean Rotation Error* (MRE). The former is defined as the Euclidean distance between the predicted and ground-truth translation vectors; the latter is computed as the geodesic distance between the predicted and reference rotation matrices.

Environments and Requirements. See Tab. 2.

Training Protocols. For this challenge, IOS-Normalizer was fine-tuned from a released checkpoint, IOS-Landmarks was used with pretrained weights, and CBCT-Landmarks was trained from scratch. The corresponding training protocols are described below and summarized in Tab. 4.

IOS-Normalizer. The model was initialized from a publicly available checkpoint pretrained on Teeth3DS [1], which includes 1,800 consistently oriented scans. It was then fine-tuned on 60 IOS from the unlabeled training set of the STS 2026 Challenge. All 60 scans were manually oriented; 50 were used for training and 10 for validation. During training, each oriented scan was downsampled to 16,384 points via FPS and either left unchanged or rotated by 180° about one of the three coordinate axes. The model was trained to identify the applied transformation, and, because a 180° rotation about an axis is self-inverse, reapplying the same transformation restores the scan to its original orientation.

IOS-Landmarks. Publicly available weights were used for both stages (segmentation and landmark prediction).

CBCT-Landmarks. We constructed a pseudo-labeled dataset from the 300 unlabeled CBCT cases in the MICCAI STS 2026 training set. Reference landmarks were obtained by transferring predictions from IOS-Landmarks to the paired CBCT volumes using an internal IOS-CBCT registration model.[3] Only 201 cases were retained after qualitative evaluation. The challenge validation set and the labeled training set (30 samples) were reserved exclusively for evaluation. Training of the CBCT landmark predictor was restricted to the *Inner* landmarks of tooth positions 1-5 in each hemi-arch, resulting in 20 output classes. Standard augmentations and foreground oversampling were applied, while left-right mirroring was disabled. The checkpoint from the fold achieving the best validation performance was used for inference. To accelerate inference, nnLandmark [7] was modified to track the maximum response and corresponding coordinate for each landmark directly across sliding-window patches, avoiding full heatmap reconstruction and CPU transfer. Batched patch processing, persistent GPU model loading, and GPU-based resampling with `resample_torch_for_nnunet` further reduced memory usage and latency.

Table 3: Ablation study on the labeled training subset and final registration performance on the validation set of the MICCAI STS 2026 Challenge.

Eval. Set	**RANSAC**	**ICP**	**MTE [mm]**	**MRE [°]**
Train	✗	✗	13.759	3.993
Train	✓	✗	8.708	2.619
Train	✗	✓	7.959	1.997
Train	✓	✓	1.813	0.891
Val.	✓	✓	4.988	1.421

4 Results and Discussion

The contribution of the individual registration components was assessed through an ablation study on the labeled training subset, for which reference IOS-CBCT

[3] Such a registration model cannot be released due to confidentiality obligations; therefore, it has not been used in the challenge itself.

Table 4: Training protocols for IOS-Normalizer and CBCT-Landmarks.

Parameter	**IOS-Normalizer**	**CBCT-Landmarks**
Batch size	1	2
Input patch size	—	128^3 voxels
Epochs	3	280
Optimizer	AdamW ($\lambda = 1e^{-4}$)	Nesterov SGD ($\mu = 0.99$)
Initial LR	$1e^{-4}$	$1e^{-2}$
LR schedule	None	Polynomial ($p = 0.9$)
Training time	4.5 min	87 h (5-fold CV)
Loss function	Cross-entropy over 4 classes	Top-20% BCE (logits)
Model parameters	13 M	102 M
FLOPs	91.3 GFLOPs/step	8.6 TFLOPs/step

transformation matrices were available. These matrices were not used to train or fine-tune any component of the proposed pipeline and were used only to quantify registration accuracy. As reported in Tab. 3, RANSAC improved registration performance both with and without ICP refinement. Point-to-plane ICP provided an additional improvement by refining the landmark-based transformation using the reconstructed surface geometry.

Quantitative Results on Validation. The proposed method achieved an MTE of 4.988 mm and an MRE of 1.421° on the challenge validation set, as reported in Tab. 3. At the submission deadline, these results placed the method first on the challenge registration leaderboard. The comparatively low rotation error indicates that the rotational component of the IOS-CBCT transformation was recovered accurately. Overall, the results support the combination of robust landmark-based initialization and local point-to-plane ICP refinement.

Qualitative Results on Validation. Qualitative inspection indicated that registration was generally more accurate when several reliable IOS-CBCT landmark correspondences were available and spatially distributed across the dental arch. Registrations based on only three or four correspondences appeared more sensitive to landmark localization errors. In these cases, ICP provided a significant improvement by refining the initial transformation using the reconstructed surface geometry. Visual examples are provided in Fig. 2.

Identification of a Systematic Transformation Labeling Error. Inspection of the 4×4 transformation matrices showed that all training cases with IDs above 269 had 3×3 linear-transformation submatrices with determinant -1, indicating a reflection across the YZ plane, as confirmed by visual comparison of the original and transformed IOS meshes. During validation, an initial submission using only positive-determinant rotations yielded an MTE of 10.037 mm and an MRE of 22.809°. Assuming the same ID-dependent inconsistency affected the validation set, we composed predictions for cases with IDs above 269 with the YZ-plane reflection; the resulting evaluation is the one previously reported in Tab. 3. The organizers later confirmed the labeling issue and our interpretation via email. Evaluation transparency was further limited by the post hoc inclusion of inference time as a weighted ranking component, despite the

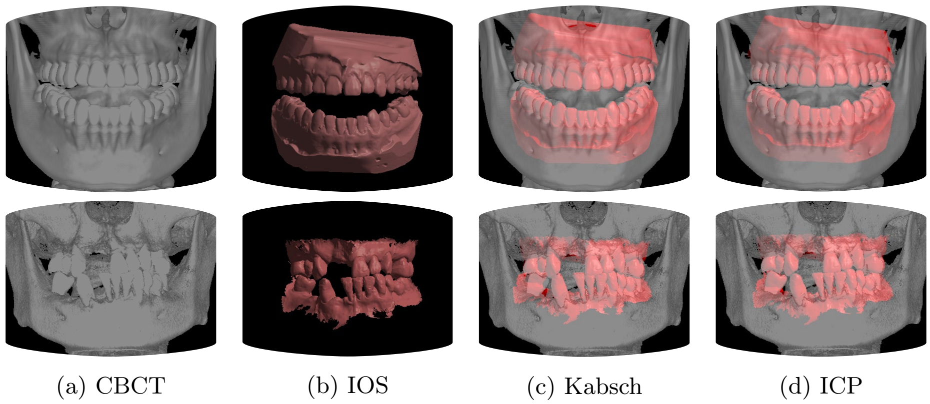

Fig. 2: Qualitative results, one patient per row (P027 and P214 from the top).

absence of a runtime metric or formal time limit in the structured submission document [21] or on the challenge website. To our knowledge, neither issue was publicly communicated, limiting reproducibility and evaluation fairness.

Limitations and Future Work. Although no visually apparent registration failures were observed on the validation set, the proposed method may fail when fewer than three geometrically consistent correspondences are available, when the selected landmarks form a nearly collinear configuration, or when localization errors place the initial alignment outside the convergence basin of ICP. Future work will investigate additional anatomical landmarks that are consistently present and can be reliably localized in both IOS and CBCT data. Increasing the number and spatial coverage of cross-modal correspondences may improve robustness in cases involving missing teeth, incomplete scans, imaging artifacts, or inaccurate landmark predictions.

5 Conclusion

This work presented a robust coarse-to-fine framework for rigid IOS-CBCT registration that combines modality-specific anatomical landmark detection, confidence-weighted RANSAC-Kabsch alignment, and point-to-plane ICP refinement. The proposed approach reduces sensitivity to large initial pose differences and inconsistent cross-modal correspondences by combining robust anatomical initialization with local surface refinement. On the challenge validation set, the method achieved an MTE of 4.988 mm and an MRE of 1.421°. These results demonstrate the effectiveness of integrating anatomical landmark correspondences with geometric surface refinement for IOS-CBCT alignment.

Acknowledgments. This work was developed and evaluated as part of the ODIN 2026 Challenge Cluster, including the STS, ToothFairy, and Bite2Text challenges [16,17,19,20].

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

References

1. Ben-Hamadou, A., Neifar, N., Rekik, A., Smaoui, O., Bouzguenda, F., Pujades, S., Boyer, E., Ladroit, E.: Teeth3DS+: An Extended Benchmark for Intraoral 3D Scans Analysis. arXiv preprint arXiv:2210.06094 (2022)
2. Besl, P., McKay, N.: A method for registration of 3-D shapes. IEEE Transactions on Pattern Analysis and Machine Intelligence **14** (1992)
3. Bolelli, F., Lumetti, L., van Nistelrooij, N., Vinayahalingam, S., Di Bartolomeo, M., Marchesini, K., Pellacani, A., Candeloro, E., Rosati, G., Xi, T., Isensee, F., Kirchhoff, Y., Krämer, L., Rokuss, M., Ulrich, C., Maier-Hein, K., Jiang, Y., Liu, Y., Wang, L., Wang, H., Chen, S., Cui, Z., Shi, P., Pan, Z., Liang, X., Ma, Q., Konukoglu, E., Wodzinski, M., Müller, H., Mai, H., Dang, X., Bhandary, S., Grosu, R., Bergé, S., Anesi, A., Grana, C.: Multi-Structure Segmentation in CBCT Volumes: the ToothFairy2 Challenge. Medical Image Analysis (2026)
4. Borghi, L., Lumetti, L., Cremonini, F., Rizzo, F., Grana, C., Lombardo, L., Bolelli, F.: Bits2Bites: Intra-oral Scans Occlusal Classification. In: International Workshop on Oral and Dental Image Analysis. pp. 45–54 (2025)
5. Charles, R.Q., Su, H., Kaichun, M., Guibas, L.J.: PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
6. Di Bartolomeo, M., Pellacani, A., Bolelli, F., Cipriano, M., Lumetti, L., Negrello, S., Allegretti, S., Minafra, P., Pollastri, F., Nocini, R., et al.: Inferior Alveolar Canal Automatic Detection with Deep Learning CNNs on CBCTs: Development of a Novel Model and Release of Open-Source Dataset and Algorithm. Applied Sciences **13**(5) (2023)
7. Ertl, A., Denner, S., Peretzke, R., Xiao, S., Zimmerer, D., Fischer, M., Bujotzek, M.R., Yang, X., Neher, P., Isensee, F., Maier-Hein, K.: nnLandmark: A Self-Configuring Method for 3D Medical Landmark Detection. In: Medical Imaging with Deep Learning (2026)
8. Fischler, M.A., Bolles, R.C.: Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. In: Fischler, M.A., Firschein, O. (eds.) Readings in Computer Vision. San Francisco (CA) (1987)
9. George, A.B., Lekshmi, P.G., Bathini, S., Govind, A.: Semi-supervised CBCT–IOS registration using PointNetLK. In: Oral and Dental Image Analysis. Lecture Notes in Computer Science, vol. 16473 (2026)
10. Jang, T.J., Yun, H.S., Hyun, C.M., Kim, J.E., Lee, S.H., Seo, J.K.: Fully automatic integration of dental CBCT images and full-arch intraoral impressions with stitching error correction via individual tooth segmentation and identification. Medical Image Analysis **93** (2024)
11. Ji, C., Liu, Y., Jiang, Y., Wang, L.: Learning-Based CBCT–IOS Registration with PointNet++ and SVD. In: Oral and Dental Image Analysis. Lecture Notes in Computer Science, vol. 16473 (2026)
12. Kabsch, W.: A solution for the best rotation to relate two sets of vectors. Acta Crystallographica Section A **32**(5) (1976)
13. Kabsch, W.: A discussion of the solution for the best rotation to relate two sets of vectors. Acta Crystallographica Section A **34**(5) (1978)
14. Kim, Y.J., Ahn, J.H., Lim, H.K., Nguyen, T.P., Jha, N., Kim, A., Yoon, J.: Novel Procedure for Automatic Registration between Cone-Beam Computed Tomography and Intraoral Scan Data Supported with 3D Segmentation. Bioengineering **10**(11) (2023)

15. Liu, J., Hao, J., Lin, H., Pan, W., Yang, J., Feng, Y., Wang, G., Li, J., Jin, Z., Zhao, Z., Liu, Z.: Deep Learning-Enabled 3D Multimodal Fusion of Cone-Beam CT and Intraoral Mesh Scans for Clinically Applicable Tooth-Bone Reconstruction. Patterns **4**(9) (2023)
16. Lumetti, L., Di Bartolomeo, M., Pellacani, A., Anesi, A., Grana, C., Bolelli, F.: Ontology-Grounded Structured Prediction for Dental CBCT Reporting. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2026 (2026)
17. Lumetti, L., Rizzo, F., Cremonini, F., Candeloro, E., Lombardo, L., Grana, C., Bolelli, F.: Do Multimodal LLMs Understand Intraoral Dental Data? Dataset, Platform, and Baselines. In: European Conference on Computer Vision (ECCV) (2026)
18. Lumetti, L., Tan, Z.Q., Borghi, L., van Nistelrooij, N., Rosati, G., Addison, O., Li, Y., Vinayahalingam, S., Grana, C., Bolelli, F.: ToothFairy3: Scaling CBCT Maxillofacial Segmentation to 77 Classes with U-Mamba2. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2026 (2026)
19. Wang, Y., Zhang, Y., Chen, X., Wang, S., Qian, D., Ye, F., Xu, F., Zhang, H., Dan, R., Zhang, Q., et al.: MICCAI 2023 STS Challenge: A retrospective study of semi-supervised approaches for teeth segmentation. Pattern Recognition **170** (2026)
20. Wang, Y., Li, Z., Wu, C., Liu, J., Zhang, Y., Ni, J., Luo, Q., Chen, J., Zhang, H., Liu, J., Han, C., Fu, K., Ji, C., Cai, X., Hao, J., Zheng, Z., Xu, S., Chen, J., Yu, X., Zhang, Q., Qian, D., Wang, S., Zhou, H.: MICCAI STS 2024 challenge: Semi-supervised instance-level tooth segmentation in panoramic X-ray and CBCT images. Medical Image Analysis **110** (2026)
21. Wang, Y., Wang, S., Li, Z.: STS 2026: The 4th Semi-supervised Teeth Segmentation Challenge on Metal Artifact Reduction and Beyond (2026)
22. Wu, X., Jiang, L., Wang, P.S., Liu, Z., Liu, X., Qiao, Y., Ouyang, W., He, T., Zhao, H.: Point Transformer V3: Simpler, Faster, Stronger. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
23. Zheng, Q., Wu, Y., Chen, J., Wang, X., Zhou, M., Li, H., Lin, J., Zhang, W., Chen, X.: Automatic multimodal registration of cone-beam computed tomography and intraoral scans: a systematic review and meta-analysis. Clinical Oral Investigations **29** (2025)