

CI-MAS: Clinically Informed Restoration-Configuration Simulation for Dental CBCT Metal Artifact Reduction

Jiatao Wang^{1*}, Zhi Li^{1*}, Jiaqi Chen¹, Zhi Han Li¹, Jinghao Sun¹, Zhengtu Wen¹, and Ping Shang¹

Hangzhou Dianzi University, Hangzhou, China
shangping@hdu.edu.cn

* These two authors contributed equally to this work.

 Corresponding author.

Abstract. Supervised metal artifact reduction (MAR) for dental cone-beam computed tomography (CBCT) depends on paired metal-free and artifact-affected images, which are rarely available clinically. Existing simulation pipelines primarily model metal geometry or physical degradation after a metal object has been specified, leaving the upstream question of where and what restoration should be simulated underexplored. We propose CI-MAS, a clinically informed simulation framework that introduces restoration configuration as an explicit decision layer before physical artifact generation. FDI-based epidemiological priors are combined with patient-specific tooth semantics to determine plausible tooth sites and restoration types, which are then instantiated as spatially aligned metal-mask volumes and processed by cone-beam forward projection, projection-domain degradation, and FDK reconstruction. Nine metal-free CBCT scans were used to generate 17 crown, filling, and mixed configurations. Under a strictly subject-disjoint train/validation/test split, retraining RISE-MAR on CI-MAS-generated data increased held-out PSNR in non-metal regions and SSIM relative to the zero-shot dental checkpoint on three unseen test subjects (PSNR 25.5 vs 23.2 dB; SSIM 0.70 vs 0.36, averaged over three training seeds). A matched tooth-agnostic random-geometry baseline achieved higher global PSNR (27.1 dB) in the current proof-of-concept setting, suggesting a potential trade-off between clinical configuration realism and geometric diversity in downstream generalization. These results support the feasibility of CI-MAS as a clinically configured CBCT artifact-simulation framework and motivate larger-scale and hybrid simulation studies.

Keywords: Metal artifact reduction · Cone-beam computed tomography · Dental CBCT · Clinical restoration prior · FDI tooth numbering · Restoration configuration · Simulation dataset.

1 Introduction

Cone-beam computed tomography (CBCT) is routinely used in oral and maxillofacial surgery, implant planning, and dental diagnosis. High-density dental

restorations interact strongly with X-rays and produce beam hardening, photon starvation, and related nonlinear effects. The reconstructed volumes therefore contain bright and dark streaks that obscure tooth boundaries and adjacent bone structures, compromising image interpretation [1, 2].

Deep learning has become a major direction in MAR research [3–7]. Supervised models, however, require paired metal-free ground truth and artifact-affected observations of the same anatomy. Such pairs are difficult to acquire clinically because dental anatomy, material composition, and scanning conditions differ before and after restorative treatment. Simulation is therefore commonly used to insert a metal object into a metal-free scan and generate a paired artifact-affected image [8].

Current simulation strategies address complementary aspects of artifact generation. Template-based approaches reuse segmented metal objects with diverse shapes and sizes to synthesize artifact-affected CT data [9], whereas procedurally generated virtual objects provide efficient geometric randomization [10]. More physics-intensive approaches, including Monte Carlo simulation, can explicitly model photon transport and material interactions at substantially higher computational cost [11]. These strategies primarily address metal representation or artifact formation after a metal object has been specified. CI-MAS instead focuses on the upstream restoration-configuration decision: which tooth is restored and which restoration morphology is instantiated before physical simulation.

In parallel, public dental datasets and automated segmentation methods have advanced rapidly. Panoramic radiograph datasets support caries segmentation and dental disease detection [12], while multi-structure cranio-maxillofacial segmentation, the MICCAI STS challenge series, and the ToothFairy line of work have improved instance-level tooth and maxillofacial structure segmentation in panoramic radiographs and CBCT [13–18]. These developments make patient-specific FDI tooth semantics available for downstream use. Nevertheless, the resulting semantics have primarily supported segmentation, diagnosis, and surgical planning rather than restoration-configuration decisions for artifact simulation.

Clinical epidemiological evidence provides an additional source of structure. Posterior teeth are the predominant sites of restorations and implants. In a retrospective analysis of 6,486 implants, posterior sites accounted for 73.6% of all implants, and mandibular first molars (FDI 36/46) were the most frequent individual sites [19]. Other population and practice-based studies likewise report a strong concentration of restorations in first permanent molars and posterior teeth [20, 21]. These observations motivate incorporating clinical tooth-site prevalence as an additional source of structure in restoration simulation.

We therefore propose CI-MAS, a clinically informed framework that connects clinical restoration priors, patient-specific FDI semantics, and projection-domain artifact simulation. The contribution of CI-MAS lies in the decision made before physical simulation: a restoration configuration specifies the target tooth site and restoration morphology, after which the configuration is instantiated as a metal-mask volume. The main contributions are:

1. **Clinical restoration-configuration modeling.** We introduce an explicit decision layer that combines FDI-based epidemiological priors with patient-specific tooth semantics to determine plausible restoration sites and types.
2. **Patient-specific configuration instantiation.** The selected semantic configuration is mapped into spatially aligned crown/filling masks while preserving the original CBCT coordinate system.
3. **Subject-disjoint downstream validation.** We evaluate CI-MAS-generated data for RISE-MAR training using a subject-disjoint train/validation/test protocol and compare it with a matched tooth-agnostic random-geometry simulation strategy.

2 Method

CI-MAS separates the clinical restoration decision from subsequent physical artifact generation. Epidemiological restoration priors and patient-specific tooth semantics first determine a restoration configuration, specifying where a restoration is placed and which morphology is instantiated. The selected configuration is then converted into a spatially aligned metal-mask volume. The metal-free CBCT and restoration mask are forward projected using the native cone-beam geometry, followed by projection-domain degradation and FDK reconstruction. Figure 1 summarizes this workflow and highlights the restoration-configuration layer introduced by CI-MAS.

2.1 Data and Tooth Semantics

All data were acquired using a Jirox dental CBCT system (Youfang Medical Technology Co., Ltd., Hefei, China). Each scan was acquired at 100 kV and 6 mA over a full 360-degree rotation with 1,600 projections and a 0.2-mm detector sampling pitch.

We screened 25 metal-free CBCT examinations and retained nine scans after checking projection completeness, assessing anatomical coverage, verifying an artifact-free baseline, and excluding severe motion, truncation, reconstruction failure, or inadequate tooth segmentation.

Raw projections were reconstructed with the vendor FDK software into volumes of $640 \times 640 \times Z$ voxels at 0.2-mm isotropic spacing. Tooth and maxillofacial structures were segmented using an nnU-Net-based pipeline [22] with ToothFairy3 weights [15], building on recent progress in the ToothFairy and STS challenge series [13, 14, 16, 17, 23] and automated cranio-maxillofacial multi-structure segmentation [18]. The resulting ToothFairy3 class labels were explicitly converted to FDI tooth identifiers before restoration instantiation. Pulp-related labels were merged with the corresponding tooth instances, and post-processing included morphological refinement and component filtering. Manual quality control subsequently verified tooth-instance completeness and anatomical correspondence between the converted FDI identifier and the segmented tooth.




Fig. 1. Overview of CI-MAS, integrating clinical tooth-site priors with patient-specific semantics for restoration configuration and metal-artifact simulation.

2.2 Clinical Restoration Configuration

Figure 2 illustrates how clinical site priors and patient-specific tooth semantics are translated into a restoration configuration and subsequently into a simulated artifact. Published epidemiological statistics inform the candidate FDI tooth sites, while crown and filling are the predefined restoration types considered in this proof-of-concept study. The patient-specific CBCT provides tooth identity and spatial extent. These inputs are combined during configuration instantiation, after which the selected restoration is converted into a spatially aligned metal-mask volume and passed to the projection-domain simulator.

Restoration prior. The candidate set included FDI 16, 17, 34, 35, 36, and 37, which represent posterior sites with comparatively high restoration or implant prevalence [19–21] under the ISO 3950 FDI designation system [24]. For this proof-of-concept study, the configuration space was intentionally restricted to these unilateral posterior sites to keep the number of simulated configurations controlled; this restriction is experimental rather than algorithmic, and the same tooth-local instantiation operators can be applied directly to contralateral sites.

Configuration instantiation. A restoration configuration represents a semantic decision rather than a metal geometry: it specifies the target FDI tooth




Fig. 2. Representative CI-MAS simulation example for a mixed restoration configuration (C37 crown and F36 filling). (a) Patient-specific tooth semantics identify the target tooth sites. (b) The selected restoration configuration is instantiated on the corresponding teeth. (c) Projection-domain simulation produces the corresponding artifact-affected CBCT. (d, e) Coronal and sagittal views visualize the resulting mask volume; the slice-wise filling approximation can appear discontinuous across planes.

site and the restoration morphology. Formally, $C = \{(t_i, r_i)\}_{i=1}^N$, where t_i denotes an FDI tooth identifier and $r_i \in \{crown, filling\}$. Given the patient-specific FDI label map L , an instantiation operator G converts the semantic configuration into a metal-mask volume, $M = G(L, C)$.

Crown. For a target tooth t , the target tooth region was dilated in each axial slice to construct a shell-like restoration mask. The 1.2-mm dilation scale was used as a literature-informed heuristic [25] reflecting the order of magnitude of restorative material space rather than a patient-specific estimate of physical crown thickness. A nominal value of 4,500 HU was assigned as a simulation parameter to represent a high-attenuation metallic restoration; this value is not intended as a quantitative material-specific CBCT measurement.

Filling. For fillings, one or two elliptical regions were sampled within the target-tooth area on each axial slice [26]. Ellipse centers were restricted to the tooth region, while semi-axis lengths and in-plane orientations were sampled within predefined ranges. The sampled ellipses were intersected with the corresponding tooth mask and assigned a nominal high-attenuation value. Because the current proof-of-concept implementation samples filling geometry slice by

slice, it does not explicitly enforce inter-slice continuity; consequently, filling masks may appear discontinuous in coronal or sagittal views.

The final 3D mask used label 1 for crowns and label 2 for fillings. Seventeen configurations were generated, covering single crowns, single fillings, dual fillings, and mixed crown-plus-filling cases; the complete configuration list is provided in Supplementary Table S1.

2.3 Projection-Domain Simulation

Forward projection and degradation. The cone-beam operator A was applied separately to the metal-free volume and metal mask, $P_{\text{clean}} = A(V_{\text{clean}})$ and $P_{\text{metal}} = A(M)$. High attenuation was modeled as $P_{\text{sim}} = P_{\text{clean}} + \alpha\beta P_{\text{metal}}$ with $\alpha = 1.35$ and $\beta = 1.12$. Beam hardening was approximated by a metal-thickness-dependent nonlinear reduction with $\gamma = 0.42$, consistent with established empirical CT metal-artifact modeling strategies [2]. A dilated metal trace was used to amplify signal-dependent noise for photon starvation with $\sigma = 0.009$ and amplification $\delta = 3.6$, and values within the trace were clipped at the 99.4th projection percentile to approximate detector saturation. These coefficients are empirical proof-of-concept parameters rather than scanner-calibrated material constants. All simulation parameters were fixed prior to downstream MAR training and remained unchanged during model training and held-out test evaluation. The degraded projections were reconstructed with FDK using the original acquisition geometry; full equations and pseudo-code are provided in Supplementary Method S3.

3 Experiments

3.1 Experimental Setup

To prevent leakage across subjects, the nine CBCT examinations were partitioned before slice extraction. Cases 2, 7, 9, 14, and 15 were used for training; case 18 was used exclusively for model selection; and cases 6, 23, and 24 were reserved for testing. This produced 724 training, 174 validation, and 392 test slices. No subject contributed slices to more than one partition. Model checkpoints were selected solely according to validation performance, and each adapted model was trained using three random seeds with random crop, flip, and rotation augmentation. The generated data were used to retrain RiseMARNet from the RISE-MAR framework [7], a four-level U-Net with CBAM attention and 64 base channels, with an L1 loss computed only in non-metal regions.

3.2 Evaluation Metrics

Peak signal-to-noise ratio (PSNR) and structural similarity (SSIM) served as the evaluation metrics. PSNR was calculated in non-metal regions ($M_{\text{bin}} = 0$) and defined as

$$\text{PSNR} = 10 \log_{10} \frac{(I_{\text{max}} - I_{\text{min}})^2}{\text{MSE}}, \quad (1)$$

where MSE is the mean squared error over non-metal pixels, $I_{\max} = 3072$ HU, and $I_{\min} = -1024$ HU. SSIM was computed slice-wise over the full slice and then averaged over metal-containing slices.

3.3 Baselines

We evaluated two comparison settings. First, the official RISE-MAR dental checkpoint (risemar_dental_preview.pkl) was used without adaptation as the zero-shot baseline. Second, we constructed a tooth-agnostic Random Geometric Mask Baseline (RGMB). RGMB masks were generated within the maxillofacial bone region using random 3D geometry while matching CI-MAS in restoration-component count, metal-volume range, subject split, training scale, projection simulator, and MAR training protocol (Supplementary Table S2). Thus, the comparison evaluates two simulation strategies rather than isolating the epidemiological tooth-site weighting alone.

4 Results

4.1 Held-out MAR Evaluation

Across the three held-out subjects, the zero-shot RISE-MAR checkpoint achieved 23.2 dB PSNR in non-metal regions and 0.36 SSIM. Retraining on CI-MAS-generated data increased the corresponding subject-level averages to 25.5 dB and 0.70 across three training seeds (values are first averaged across configurations within each held-out subject, then across seeds). The gain was observed under the subject-disjoint train/validation/test split, eliminating the slice-level leakage ambiguity of the original evaluation.

4.2 Simulation-Strategy Comparison

Under the matched training protocol, RGMB achieved a higher subject-level global PSNR of 27.1 dB than CI-MAS. Thus, the present results do not support a claim that clinically informed configuration is universally superior for downstream MAR performance. Instead, they indicate that broader geometric randomization can provide strong global generalization even without tooth-level restoration semantics. Exploratory regional evaluation suggested different behavior near the simulated metal: CI-MAS obtained higher metal-adjacent PSNR in two of three training seeds, whereas RGMB remained stronger in regions farther from metal. Because this pattern was not consistent across all seeds, we treat it as hypothesis-generating evidence rather than a demonstrated regional advantage.

4.3 Exploratory Regional Analysis

Figure 4 reports an exploratory spatial analysis of the held-out PSNR as a function of distance from the simulated metal, complementing the subject-level comparison in Fig. 3.




Fig. 3. Held-out MAR performance under the strictly subject-disjoint protocol. Bars show subject-level means, computed by first averaging configurations within each held-out subject (case6, case23, case24); error bars denote the SD across the three subjects. CI-MAS and RGMB results are averaged over three training seeds; the SD of the subject-level mean across seeds was 0.87 dB (CI-MAS) and 0.16 dB (RGMB).

5 Conclusion

CI-MAS introduces restoration configuration as an explicit clinically informed decision layer before physical CBCT artifact simulation. By combining epidemiological tooth-site information with patient-specific FDI semantics, the framework generates restoration configurations that can be instantiated into aligned metal-mask volumes and used for projection-domain MAR simulation. Under a subject-disjoint evaluation, CI-MAS-generated data improved RISE-MAR over zero-shot inference on unseen subjects, supporting the feasibility of clinically configured simulation for downstream MAR training.

The matched RGMB comparison achieved higher global PSNR, indicating that clinical configuration realism and geometric diversity are distinct simulation objectives rather than a simple performance hierarchy. Limitations include the small single-scanner cohort, restricted unilateral tooth-site set, empirically parameterized degradation model, and slice-wise filling approximation. The tooth-local operators also extend to contralateral sites. Future work should combine clinically informed placement with broader volumetric geometry, multi-center data, and physics-validated evaluation.

Data and Reproducibility Statement. The nine CBCT scans were selected from the publicly released CSS 2026 CBCT Metal Artifact Removal Challenge dataset (<https://tianchi.aliyun.com/competition/entrance/532440>) and were collected with a Jirox dental CBCT system. No new participant recruitment or prospective data collection was performed; analyses used de-identified challenge data under the released data-use terms. The projection-domain simulation pipeline is described in pseudo-code in Supplementary Method S3, and the implementation will be made publicly available upon publication.




Fig. 4. Metal-artifact removal by distance from the metal: CI-MAS vs. RGMB. PSNR is computed in non-metal pixels binned by distance from the nearest metal voxel (2-mm bins), averaged over the five held-out configurations; error bars denote the SD across configurations. Profiles are averaged over three training seeds. CI-MAS is higher within 2 mm of the metal, whereas RGMB is higher beyond 2 mm. Given the limited sample, this is treated as hypothesis-generating rather than a demonstrated regional advantage.

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

References

1. Schulze, R., Heil, U., Gross, D., Bruellmann, D.D., Dranischnikow, E., Schoemer, E.: Artefacts in CBCT: a review. *Dentomaxillofacial Radiology* 40(5), 265–273 (2011).
2. De Man, B., Nuyts, J., Dupont, P., Marchal, G., Suetens, P.: Metal streak artifacts in X-ray computed tomography: a simulation study. *IEEE Trans. Nucl. Sci.* 46(3), 691–696 (1999).
3. Zhang, Y., Yu, H.: Convolutional neural network based metal artifact reduction in X-ray computed tomography. *IEEE Trans. Med. Imaging* 37(6), 1370–1381 (2018).
4. Lin, W., et al.: DuDoNet: dual domain network for CT metal artifact reduction. In: *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pp. 10504–10513 (2019).
5. Wang, H., et al.: InDuDoNet+: a deep unfolding dual domain network for metal artifact reduction in CT images. *Med. Image Anal.* 85, 102729 (2023).
6. Luo, M., et al.: Bi-Constraints Diffusion: a conditional diffusion model with degradation guidance for metal artifact reduction. *IEEE Trans. Med. Imaging* 44(9), 3552–3562 (2025).
7. Ma, C., et al.: RISE-MAR: radiologist-in-the-loop self-training for generalizable CT metal artifact reduction. *IEEE Trans. Med. Imaging* 44(6), 2504–2514 (2025).
8. Agrawal, H., et al.: Deep learning based projection domain metal segmentation for metal artifact reduction in cone beam computed tomography. *IEEE Access* 11, 100371–100382 (2023).

9. Yu, L., Zhang, Z., Li, X., Xing, L.: Deep Sinogram Completion with Image Prior for Metal Artifact Reduction in CT Images. *IEEE Trans. Med. Imaging* 40(1), 228–238 (2021).
10. Haneda, E., Peters, N., Zhang, J., et al.: AAPM CT Metal Artifact Reduction Grand Challenge. *Med. Phys.* 52(10), e70050 (2025).
11. Nasirudin, R.A., Mei, K., Penchev, P., et al.: Reduction of Metal Artifact in Single Photon-Counting Computed Tomography by Spectral-Driven Iterative Reconstruction Technique. *PLOS ONE* 10(5), e0123074 (2015).
12. Zhang, Y., Ye, F., Chen, L., et al.: Children’s dental panoramic radiographs dataset for caries segmentation and dental disease detection. *Sci. Data* 10(1), 380 (2023).
13. Bolelli, F., et al.: Segmenting the inferior alveolar canal in CBCT volumes: the ToothFairy challenge. *IEEE Trans. Med. Imaging* (2025). doi: 10.1109/TMI.2024.3523096.
14. Wang, Y., Zhang, Y., Chen, X., et al.: MICCAI 2023 STS Challenge: a retrospective study of semi-supervised approaches for teeth segmentation. *Pattern Recognit.* 170, 112049 (2026).
15. Lumetti, L., Tan, Z.Q., Borghi, L., et al.: ToothFairy3: scaling CBCT maxillofacial segmentation to 77 classes with U-Mamba2. In: 29th International Conference on Medical Image Computing and Computer-Assisted Intervention (2026).
16. Bolelli, F., Marchesini, K., van Nistelrooij, N., et al.: Segmenting maxillofacial structures in CBCT volume. In: *IEEE/CVF Conf. Comput. Vis. Pattern Recognit.* (2025).
17. Lumetti, L., Pipoli, V., Bolelli, F., Ficarra, E., Grana, C.: Enhancing patch-based learning for the segmentation of the mandibular canal. *IEEE Access* 12, 79014–79024 (2024).
18. Bao, J., Tan, Z., Sun, Y., et al.: Deep ensemble learning-driven fully automated multi-structure segmentation for precision craniomaxillofacial surgery. *Frontiers in Bioengineering and Biotechnology* 13, 1580502 (2025).
19. Huermuezlue, M.K., Mollaoğlu, N.: Dental implant procedures performed in Gazi University Faculty of Dentistry: prevalence and demographic distribution. *Cumhuriyet Dental Journal* 27(4), 240–244 (2024).
20. Beltran-Aguilar, E.D., et al.: Prevalence and mean number of teeth with amalgam and nonamalgam restorations, United States, 2015 through 2018. *Journal of the American Dental Association* 154(5), 417–426 (2023).
21. da Silva Pereira, D., et al.: Practice-based analysis of direct posterior dental restorations performed in a public health service: retrospective long-term survival in Brazil. *PLOS ONE* 15(12), e0243288 (2020).
22. Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* 18, 203–211 (2021).
23. Wang, Y., Li, Z., Wu, C., et al.: MICCAI STS 2024 Challenge: semi-supervised instance-level tooth segmentation in panoramic X-ray and CBCT images. *arXiv:2511.22911* (2025).
24. ISO 3950:2016, Dentistry — Designation system for teeth and areas of the oral cavity. International Organization for Standardization, Geneva (2016).
25. Ram, H.K., Shah, R.J., Agrawal, H.S.: Evaluation of three different tooth preparation techniques for metal ceramic crowns by comparing preparation depths: an in vitro study. *J. Indian Prosthodont. Soc.* 15(4), 314–320 (2015).
26. Jokstad, A.: The dimensions of everyday class-II cavity preparations for amalgam. *Acta Odontol. Scand.* 47(2), 89–98 (1989).