CrownFusion: 3D Dental Crown Generation Using Geometry Images and Latent Diffusion

Johan Ziruo Ye[1,2*], Xingguang Yan[3*] (✉), Søren Hauberg[1], Angel X. Chang[3,4], Hao Zhang[3], and Peter Lempel Søndergaard[2]

[1] Technical University of Denmark, Kongens Lyngby, Denmark
[2] 3Shape, Copenhagen, Denmark
[3] Simon Fraser University, Vancouver, Canada
[4] Canada-CIFAR AI Chair, Amii
qheldiv@gmail.com

Abstract. We introduce *CrownFusion*[†], a latent diffusion model to generate 3D dental crown designs from *geometry images*, a 2D-grid encoding of surface coordinates. Departing from traditional approaches that use point clouds, our "*CrownImage*" representation not only unlocks the full power of image-based generative architectures, but also preserves fine geometric detail by operating at far higher spatial resolutions than point clouds to alleviate smoothing artifacts in the 3D generation. Conditioned on jointly learned embeddings of the surrounding dentition, our probabilistic diffusion model produces anatomically plausible crowns with morphological variations tailored to individual patient cases. To accommodate multiple valid designs, we propose an occlusal fit metric based on proximity to the antagonist teeth. We evaluate our method through quantitative experiments and qualitative studies involving expert testimonials. We also release the *FDI 16 Crown Dataset*, comprising 11,513 annotated intra-oral scans, to support future research in dental CAD.

Keywords: Dental Crowns · Shape Generation · 3D Latent Diffusion

1 Introduction

Computer-aided dental design software has streamlined the creation of artificial tooth replacements, allowing technicians to adapt templates to individual patients. Yet, crown design remains labor-intensive, requiring precise adaptation to each patient's unique dental morphology. A clinically acceptable crown must achieve adequate occlusal fit with the opposing dentition, maintain functional contact without hyperocclusion [3], and reflect natural tooth morphology, including well-defined occlusal grooves that direct food flow during mastication. In practice however, there is no single correct crown for a given case. Multiple clinically acceptable designs may exist, shaped by the technician's training, experience, and aesthetic judgement [13, 19]. This motivates a generative approach to

* Equal contribution.
† Code and data available at `https://github.com/JohanYe/CrownFusion`.

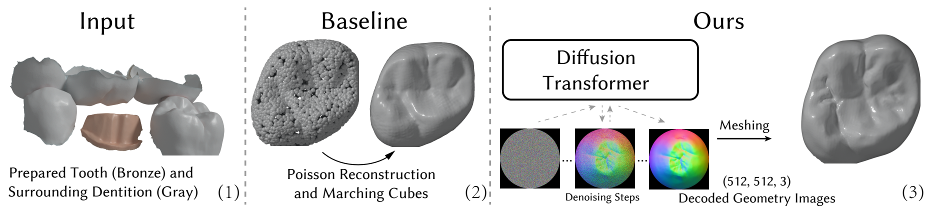

Fig. 1. Overview. From left to right: (1) Input: prepared tooth (bronze) and surrounding dentition (gray) from an intraoral scan. (2) Baseline: point cloud-based methods with Poisson reconstruction yield overly smooth geometry eroding occlusal detail. (3) Ours: a diffusion transformer generates a geometry image encoding of the crown surface, which is then meshed to recover the 3D crown, preserving fine occlusal features.

capture a *space* of clinically acceptable solutions rather than predicting a single outcome, as well as developing a suitable metric for evaluation.

In terms of shape representations, previous approaches to automated dental crown design face fundamental limitations. Point cloud-based methods [7, 20, 23] frame crown design as a shape completion task, requiring a separate mesh extraction step that typically introduces surface artifacts. They tend to produce overly smooth outputs and erode clinically important details such as sharp fissures, due to the difficulty of placing points at high-curvature features and optimization using the Chamfer Distance [1] (CD), which is insensitive to local geometric details. Depth image-based methods [8, 19, 27] are confined to projection from a single viewpoint, inherently discarding 3D information. For instance, the lingual surface of an antagonist tooth, which directly contacts the generated crown during occlusion, may be entirely occluded and absent from the depth rendering.

In this work, we address these limitations by generating 3D dental crowns through *geometry images* [6], which encode a tooth's surface geometry into a 2D grid, i.e., a *CrownImage*, with each entry storing an $[x, y, z]$ coordinate on the surface as color (Fig. 2). Such a 3D-to-2D mapping is facilitated by the reliable manifold topology of tooth meshes, ensuring a seamless parameterization onto the 2D grid. Unlike single-view depth renderings, geometry images retain full 3D geometric information, while supporting far higher spatial resolutions than point clouds to better preserve fine-grained and high-curvature teeth features.

CrownImages, our image-based teeth representation, enables us to develop *CrownFusion*, a conditional latent diffusion model for 3D crown generation. To our knowledge, this is both the first use of geometry images and the first latent diffusion model for this task. We combine the Scalable Interpolant Transformer (SiT) [11] and a FLUX VAE [2], with cross-attention conditioning on shape embeddings of the surrounding dentition (Fig. 1). The generation of multiple morphologically varied crown candidates for a single patient offers clinicians a range of viable designs to explore. To complement standard CD evaluation, we propose an occlusal fit metric based on proximity to the antagonist teeth to better capture functional performance. Experiments show that CrownFusion achieves competitive occlusal alignment with existing completion-based methods

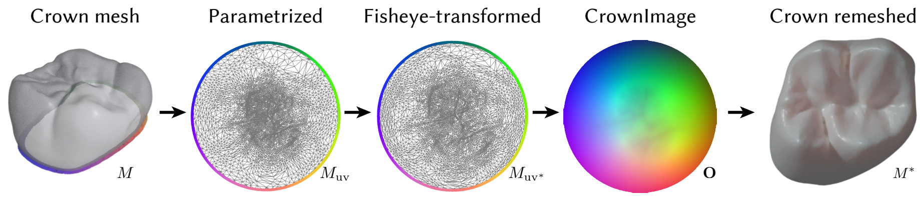

Fig. 2. CrownImage construction pipeline. M: crown mesh. M_{uv}: parametrized mesh on the unit disk via boundary-fixed harmonic mapping. $M_{\mathrm{uv}*}$: fisheye-transformed UV mesh (inner disk expanded, outer annulus compressed). $\mathbf{O}$: rasterized geometry image encoding (x, y, z) and α. M^*: remeshed crown recovered by triangulating $\mathbf{O}$.

while producing substantially fewer intersections. Expert testimonials confirm that our method consistently produces more well-defined and natural-looking occlusal anatomy and was less prone to blunt or anatomically incoherent surfaces than baselines. Finally, we release the *FDI 16 Crown Dataset* of 11,513 annotated intraoral scans.

2 Related Work

Crown design via 3D completion. Recent methods frame dental crown generation as a 3D shape completion task. *DCrownFormer* [23] and *DCrownFormer+* [24] encode antagonist and preparation tooth point clouds via transformer blocks and a morphology-aware cross-attention decoder. Similarly, *DMC* [7] designs a transformer-based encoder-decoder followed by FoldingNet [25] to output a point cloud. Recently, *VBCD* [20] voxelizes intraoral scans and generates coarse crowns via a 3D UNet followed by point cloud refinement. Relying on Shape-as-Points [14] for final mesh extraction, these methods are all deterministic completion approaches with a single outcome per input. In contrast, CrownFusion is fully generative and produces multiple designs based on the surrounding teeth.

Crown design via depth image synthesis. Converting dental meshes to depth images is another way to enable the use of 2D deep learning methods. Hwang et al. [8] applied GANs [5] to synthesize crown depth images, though results varied in anatomical correctness. Yuan et al. [27] extended this with a conditional GAN incorporating neighboring teeth and an occlusal groove filter loss to improve fissure alignment. Tian et al. [19] proposed a dual-GAN framework with global and local discriminators to enforce both contextual consistency and morphological detail. Like CrownFusion, these methods are generative, but they lack the controllable generation afforded by diffusion models. In addition, generating a complete crown represented by depth images is challenging as it may require generating multiple views that are consistent.

Geometry images. Geometry images have primarily been explored for 3D objects with simple topology. Sinha et al. [18] first applied 2D deep learning to geometry images for shape classification, using authalic spherical parameterization

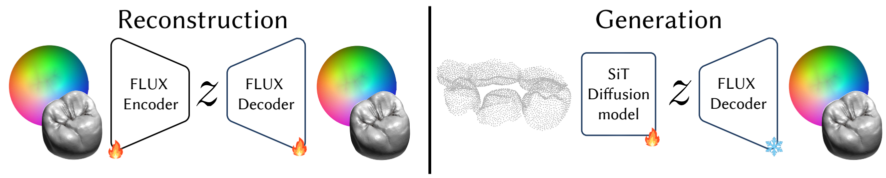

Fig. 3. Overview of CrownFusion. Left: VAE reconstruction from geometry images. Right: conditional generation via SiT diffusion in the learned latent space.

to unwrap closed meshes. Maron et al. [12] extended this to toric objects with an additional dimension and stitching to handle topological mismatches. Recently, Yan et al. [22] proposed a latent diffusion pipeline that segments objects into simple patches, maps each to geometry images, and stitches them during reconstruction. In contrast, intraoral scans are open surfaces, inherently compatible with geometry images without complex parameterization or stitching.

3 Approach

Our core contribution is the introduction of CrownImages, geometry images [6] of dental crowns, as a data representation that enables the use of established 2D generative architectures. We train CrownFusion, a latent diffusion model that generates CrownImages conditioned on the surrounding dentition, and contribute the FDI 16 Crown Dataset for training and evaluation. Next, we describe the parameterization, our diffusion pipeline, and the dataset.

CrownImage construction. We represent the dental crown mesh M as a 3-channel geometry image $\mathbf{O} \in \mathbb{R}^{H\times W\times 3}$, with the 3 channels encoding the spatial coordinates (x, y, z). Figure 2 shows the process via which we obtain the CrownImage $\mathbf{O}$ for a given mesh M. Because a tooth is a topological zero-genus surface, we use a global parameterization to encode the tooth with a single chart [17]. To construct $\mathbf{O}$ for a given mesh M, we first flatten M into a disk M_{uv} via boundary-fixed harmonic parameterization [4]. We then apply a fisheye-style radial transform to obtain $M_{\mathrm{uv}*}$ so that more space is allocated for the regions closer to the center (corresponding to occlusal surfaces) and less space for the outer rim (containing the axial surfaces). To obtain $\mathbf{O}$, we rasterize the inverse position map $\pi^{-1} : \mathcal{S} \to \mathbb{R}^3$ onto a regular $H \times W \times 3$ grid. The mesh M^* is recovered by triangulating adjacent pixels [22].

Latent diffusion pipeline. We adopt a latent diffusion framework [15] to operate on CrownImages (Fig. 3). A variational autoencoder (VAE) encodes the images into latent codes, and a diffusion model generates dental crowns conditioned on the surrounding dentition.

For the VAE, we finetune FLUX [2]. Unlike natural images, CrownImages encode continuous spatial coordinates rather than color values, resulting in a distribution that differs substantially from the data FLUX was pretrained on. Moreover, small latent space errors correspond directly to geometric distortions on the reconstructed surface, making reconstruction fidelity critical. Finetuning the VAE is therefore necessary to faithfully capture geometric variability.

For the diffusion model, we train a Scalable Interpolant Transformer (SiT) [11] from scratch. The model is conditioned on two complementary inputs: the prepared tooth's neighbors, the antagonist, and the antagonist's neighbors, and the margin line of the prepared tooth. The margin line delineates the preparation boundary, anchoring the crown's seating position and constraining the generated geometry from below. To retain 3D spatial information, we represent the surrounding dentition as a point cloud sampled from the prepared tooth's neighbors and its antagonist. This context is encoded using a shape encoder [29], trained jointly with the diffusion model. The resulting embedding is integrated via cross-attention at each denoising step, providing spatial and morphological cues from surrounding teeth. This enables the model to infer occlusal surface features from the antagonist and cusp and fissure placement from the neighboring teeth. The decoded CrownImage is meshed by triangulating adjacent occupied pixels, i.e. pixels with valid geometric coordinates [22].

FDI 16 crown dataset. To advance dental crown generation research, we release the *FDI 16 crown dataset*. The dataset comprises 11,513 anonymized intraoral scans collected through the 3Shape Dental System, with 1,000 scans held out for validation and testing respectively. Each scan is paired with the final crown restoration. We focus on the FDI 16 tooth (upper right first molar) as it is a frequently restored tooth and presents the most complex occlusal morphology of any tooth type, with four to five cusps, an oblique ridge, and a detailed fissure pattern, making it a challenging benchmark. The FDI 16 Tooth Dataset [26] also includes this tooth, but it only contains isolated natural teeth. Our dataset provides crown restorations paired with surrounding dentition, enabling conditional generation. All scans are centered on the FDI 16 crown with the gingiva removed via segmentation using 3Shape Ortho System 2025. Where available, the scans include adjacent teeth on the preparation side (FDI 15, 17) and opposing teeth on the antagonist side (FDI 45, 46, 47). We exclude crowns with low overall curvature or significant antagonist intersections, retaining only geometrically complex, well-fitting samples. The crown restorations were designed by dental technicians in routine clinical workflows, reflecting real-world variability in crown quality.

4 Experiments

We evaluate the reconstruction fidelity of CrownFusion on the FDI 16 Tooth Dataset [26], and the generative quality and diversity of the full pipeline on our FDI 16 Crown Dataset through standard metrics and a proposed measure occlusal fit metric to complement them.

Autoencoding using geometry images and comparisons. Since the quality of crown generation via CrownFusion is bounded by the reconstruction fidelity of its autoencoder (AE), we first evaluate our VAE to establish an upper bound on generative performance. In the absence of prior image-based AE

Table 1. We report reconstruction errors, in CD and EMD, for the FLUX-based VAE in our CrownFusion and several prior point cloud-based baselines.

	DPM [10]	SetVAE [9]	LION [28]	FN [25]	VF-Net [26]	CrownFusion†
CD ($\times 10^2$)	10.04	21.50	5.35	5.26	1.21	**0.02**
EMD ($\times 10^2$)	43.98	59.24	22.85	33.67	6.30	**0.45**

baselines for our task, we compare against point cloud-based autoencoders previously benchmarked on the FDI 16 Tooth Dataset [26], reporting CD [1] and Earth mover's distances (EMDs) [16]. It is worth noting that these point cloud AE baselines all input and output 2,048 points, while our VAE inputs and outputs geometry images. We convert the output CrownImage into a mesh, project each GT point onto its surface, and compute CD and EMD between the GT and projected point clouds (denoted CrownFusion†). This point-to-surface evaluation is more favorable than point-to-point comparison, as the continuous mesh surface naturally reduces distances to arbitrary query points.

Table 1 shows that our VAE yields the lowest error, which can be attributed to the significantly higher spatial resolution of the CrownImages and the FLUX architecture. Visual comparison in Fig. 4 shows that our reconstructions preserve geometric detail beyond what the point-based protocol could capture.

Static occlusal fit. Next, we train and evaluate the conditional generative pipeline on FDI 16. We compare against DMC [7] and VBCD [20], the two most recent crown completion-based methods with available implementations (Table 2). We report the Chamfer distance to a dental crown selected by a dental technician, following prior work [7, 20, 23]. Although our generative approach yields higher reconstruction error than the deterministic VBCD, it outperforms DMC, despite modeling a full distribution rather than a single output.

Relying on a single technician-designed crown as ground truth is problematic, since crown designs vary substantially between technicians, and many retain most of the default library tooth morphology with minor adjustments. Similarity to one such reference therefore says little about the clinical quality of a generated crown [13, 21]. We propose to evaluate occlusal fit by measuring the proximity between the generated occlusal surface and the antagonist dentition. Since teeth rarely align one-to-one, we evaluate occlusal fit against the antagonist tooth and its two nearest neighbors. We sample 30,000 points from these antagonist teeth

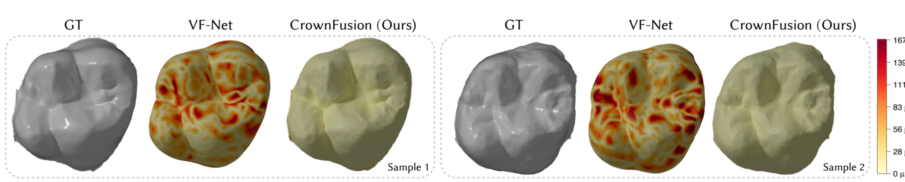

Fig. 4. Surface distance heatmaps to GT show VF-Net's errors concentrated at cusps and fissures, while our geometry image VAE is close to GT across the crown surface.

Table 2. Comparison of CrownFusion against previous methods for dental crown design. Crown similarity is measured by Chamfer Distance (CD-L2, mm) to a dental technician selected crown. Occlusal fit is evaluated against the three nearest antagonist teeth (by mean distance of 10th-percentile closest points) and surface intersections (count out of 1,000 sampled points, and area in mm^2). CrownFusion achieves the best occlusal fit with fewer intersections and competitive antagonist proximity.

	Crown Similarity ↓	Bite Proximity ↓	Intersections ↓	
Model	CD-L2 (mm)	10% closest	num	area (mm^2)
DMC [7]	0.433	0.734	256	1.110
VBCD [20]	**0.353**	0.747	143	2.363
CrownFusion (ours)	0.412	**0.733**	**127**	**0.612**
GT		0.716	12	0.0437

to enable dense evaluation, and report the mean distance of the closest 10th-percentile of antagonist points to the generated crown surface. This percentile captures the near-contact region most relevant to the occlusal fit. We also report intersections between the generated crown and antagonist surfaces, both count and area, as occlusal intersections require significant manual sculpting to correct, unlike proximal intersections which occur on smoother, more accessible surfaces. Minimizing these intersections is therefore essential for reducing manual intervention in clinical workflows.

CrownFusion achieves the closest occlusal alignment to the antagonist teeth (Table 2), with fewer intersections than both baselines and roughly half the intersection rate of DMC. When intersections do occur, their area is substantially smaller than in prior methods. Figure 5 shows a representative example. Ground truth crowns exhibit minimal intersections (12/1000, $0.0437\,\mathrm{mm}^2$), indicating room for future improvement.

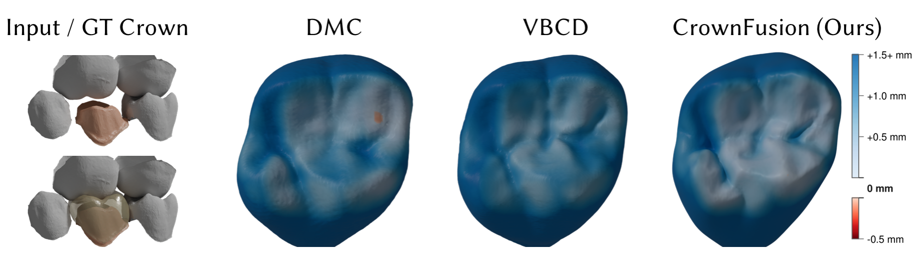

Fig. 5. Generation comparison. **Left:** dental arch with prepared abutment (top) and the reference crown in context (bottom). **Right:** occlusal views of crowns generated by DMC, VBCD, and CrownFusion, colored by signed distance to the antagonist (blue: gap; red: intersection). VBCD shows a large uniform gap, indicating poor occlusal contact. DMC exhibits a localized intersection (red). CrownFusion achieves the closest antagonist proximity with no intersections, and preserves well-defined occlusal anatomy.

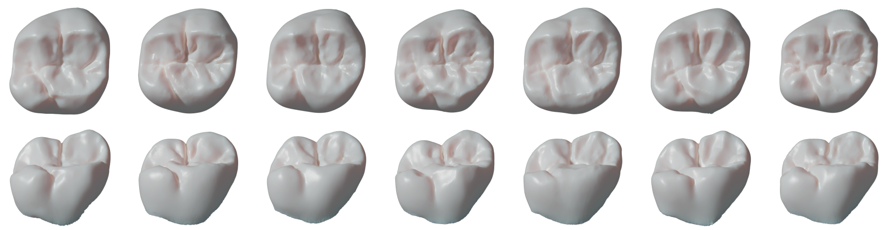

Fig. 6. Generated crown samples for the same input. CrownFusion captures expected axes of variation: fissure depth, secondary fossae detail, and distolingual cusp morphology vary, while marginal ridge height and overall crown outline remain largely stable.

Expert evaluation. To complement quantitative metrics, we consulted two dental professionals who assessed generated crowns across all methods. They noted that CrownFusion produces well-defined occlusal anatomy, rarely generating the featureless surfaces typical of DMC's overly smooth crowns with shallow fissures and blunt cusps. Across DMC, VBCD, and even the reference crowns, fissures occasionally appear as isolated incisions without corresponding cusp development, yielding anatomically incoherent surfaces; CrownFusion was less prone to this. It does, however, occasionally produce flattened proximal surfaces that conform too closely to adjacent teeth rather than forming convex contacts, likely reflecting training data where technicians rely on automated fitting tools. This, with occlusal anatomy that does not always match the clinical context, points to conditioning and data curation as areas for improvement.

Morphological variation across samples. Finally, we examine whether the variation across samples generated for the same patient aligns with the design freedom available in clinical practice. While marginal ridge height and gross cusp placement are constrained by the surrounding dentition, the occlusal surface is the primary locus of design freedom [21], with cusp height, fissure depth, and secondary ridge anatomy varying substantially between technicians.

Fig. 6 shows how the model varies fissure pattern detail, secondary fossae, distolingual cusp, and oblique ridge morphology across samples, consistent with these features carrying more design freedom. Marginal ridge height and overall crown outline remain largely stable, though in isolated cases the variation is larger than clinically ideal. Notably, the features that vary are those most difficult to sculpt manually, as existing CAD tools primarily support local adjustments rather than global morphological changes. Extending the model to offer controllable gross shape diversity through explicit conditioning is a natural next step.

Limitations. Our evaluation focuses on FDI 16; generalization across all tooth types remains to be validated, though the strong generalization of latent diffusion in image synthesis is encouraging. The disk topology of geometry images currently restricts the framework to single-tooth restoration, and conditioning on absent neighbouring teeth is not yet addressed. Generation time ($\sim$5 min per 8 samples on an RTX 4090) is a practical constraint, though the extensive literature on diffusion sampling acceleration transfers directly to our setting.

5 Conclusion

We presented *CrownFusion*, the first latent diffusion model for 3D dental crown generation using an image-based teeth representation, *CrownImage*. CrownFusion generates diverse, well-formed crowns with improved handling of high-curvature dental features, competitive occlusal fit, and fewer antagonist intersections. We release an accompanying dataset of 11,513 annotated intraoral scans and propose an occlusal fit metric based on antagonist proximity.

Acknowledgments. This work was partly funded by Innovation Fund Denmark (1044-00172B). SH was supported by the Novo Nordisk Foundation through the Center for Basic Machine Learning Research in Life Science (NNF20OC0062606), VILLUM FONDEN (42062), and the European Research Council (ERC) under the European Union's Horizon programme (grant agreement 101125993). AXC is supported by a CIFAR AI Chair and a NSERC Discovery Grant. HZ is supported by a NSERC Discovery Grant.

Disclosure of Interests. J. Z. Ye and P. L. Søndergaard are employees of 3Shape A/S, and the dataset used in this work was collected through 3Shape software. The remaining authors have no competing interests to declare that are relevant to the content of this article.

Bibliography

[1] Barrow, H.G., Tenenbaum, J.M., Bolles, R.C., Wolf, H.C.: Parametric correspondence and chamfer matching: Two new techniques for image matching. In: IJCAI, pp. 659–663 (1977)

[2] Black-Forest-Labs: FLUX (2024), https://github.com/black-forest-labs/flux

[3] Bronson, M.R., Lindquist, T.J., Dawson, D.V.: Clinical acceptability of crown margins versus marginal gaps as determined by pre-doctoral students and prosthodontists. J. Prosthodont. **14**(4), 226–232 (2005)

[4] Eck, M., DeRose, T., Duchamp, T., Hoppe, H., Lounsbery, M., Stuetzle, W.: Multiresolution analysis of arbitrary meshes. In: SIGGRAPH '95, pp. 173–182 (1995)

[5] Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. In: NeurIPS (2014)

[6] Gu, X., Gortler, S.J., Hoppe, H.: Geometry images. ACM Trans. Graph. **21**(3), 355–361 (2002)

[7] Hosseinimanesh, G., Ghadiri, F., Guibault, F., Cheriet, F., Keren, J.: From mesh completion to AI designed crown (2025), arXiv:2501.04914

[8] Hwang, J.J., Azernikov, S., Efros, A.A., Yu, S.X.: Learning beyond human expertise with generative models for dental restorations (2018), arXiv:1804.00064

[9] Kim, J., Yoo, J., Lee, J., Hong, S.: SetVAE: Learning hierarchical composition for generative modeling of set-structured data. In: CVPR (2021)

[10] Luo, S., Hu, W.: Diffusion probabilistic models for 3D point cloud generation. In: CVPR (2021)

[11] Ma, N., Goldstein, M., Albergo, M.S., Boffi, N.M., Vanden-Eijnden, E., Xie, S.: SiT: Exploring flow and diffusion-based generative models with scalable interpolant transformers. In: ECCV (2024)

[12] Maron, H., Galun, M., Aigerman, N., Trope, M., Dym, N., Yumer, E., Kim, V.G., Lipman, Y.: Convolutional neural networks on surfaces via seamless toric covers. ACM Trans. Graph. **36**(4), 71:1–71:10 (2017)

[13] McCracken, M.S., Litaker, M.S., Thomson, A.E.S., Slootsky, A., Gilbert, G.H.: Laboratory technician assessment of the quality of single-unit crown preparations and impressions as predictors of the clinical acceptability of crowns. J. Prosthodont. **29**(2), 114–123 (2020)

[14] Peng, S., Jiang, C.M., Liao, Y., Niemeyer, M., Pollefeys, M., Geiger, A.: Shape as points: A differentiable poisson solver. In: NeurIPS (2021)

[15] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022)

[16] Rubner, Y., Tomasi, C., Guibas, L.J.: The earth mover's distance as a metric for image retrieval. Int. J. Comput. Vis. **40**(2), 99–121 (2000)

[17] Sander, P.V., Wood, Z.J., Gortler, S., Snyder, J., Hoppe, H.: Multi-chart geometry images (2003)
[18] Sinha, A., Bai, J., Ramani, K.: Deep learning 3D shape surfaces using geometry images. In: ECCV, pp. 223–240 (2016)
[19] Tian, S., Huang, R., Li, Z., Fiorenza, L., Dai, N., Sun, Y., Ma, H.: A dual discriminator adversarial learning approach for dental occlusal surface reconstruction. J. Healthc. Eng. **2022**, 1933617 (2022)
[20] Wei, L., Liu, C., Zhang, W., Zhang, Z., Zhang, S., Li, H.: VBCD: A voxel-based framework for personalized dental crown design (2025), arXiv:2507.17205
[21] Xie, B.y., He, X., Hu, L., Guo, S.l., Chen, J.l., Zhang, J., Shen, X.q., Geng, Y.m., Li, W.: Morphological comparison between artificial intelligence-driven and manual CAD design in single tooth restoration. BMC Oral Health **25**(1), 1633 (2025)
[22] Yan, X., Lee, H.H., Wan, Z., Chang, A.X.: An object is worth 64x64 pixels: Generating 3D object via image diffusion (2024), arXiv:2408.03178
[23] Yang, S., Han, J., Lim, S.H., Yoo, J.Y., Kim, S., Song, D., Kim, S., Kim, J.M., Yi, W.J.: DCrownFormer: Morphology-aware point-to-mesh generation transformer for dental crown prosthesis from 3D scan data of antagonist and preparation teeth. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024, pp. 109–119 (2024)
[24] Yang, S., Han, J.Y., Lim, S.H., Kim, S., Lee, J., Kim, K.S., Kim, J.M., Yi, W.J.: DCrownFormer+: Morphology-aware mesh generation and refinement transformer for dental crown prosthesis. Med. Image Anal. **105**, 103717 (2025)
[25] Yang, Y., Feng, C., Shen, Y., Tian, D.: FoldingNet: Point cloud auto-encoder via deep grid deformation. In: CVPR, pp. 206–215 (2018)
[26] Ye, J.Z., Ørkild, T., Søndergaard, P.L., Hauberg, S.: Variational autoencoding of dental point clouds (2023), arXiv:2307.10895
[27] Yuan, F., Dai, N., Tian, S., Zhang, B., Sun, Y., Yu, Q., Liu, H.: Personalized design technique for the dental occlusal surface based on conditional generative adversarial networks. Int. J. Numer. Methods Biomed. Eng. **36**(5), e3321 (2020)
[28] Zeng, X., Vahdat, A., Williams, F., Gojcic, Z., Litany, O., Fidler, S., Kreis, K.: LION: Latent point diffusion models for 3D shape generation. In: NeurIPS (2022)
[29] Zhang, B., Tang, J., Niessner, M., Wonka, P.: 3DShape2VecSet: A 3D shape representation for neural fields and generative diffusion models. ACM Trans. Graph. **42**(4) (2023)