

FDMMBone: Flow-Based Conditional Deformable Mesh Modeling for Accurate 3D Bone Defect Completion

Zhenhong Liu¹, Xingce Wang^{1✉}, Zhongke Wu¹, Xiaodong Ju², Jingyi Liu³, Yi-Cheng Zhu^{3✉}, Wuyang Shui⁴

¹ School of Artificial Intelligence, Beijing Normal University, Beijing, China
wangxingce@bnu.edu.cn

² Department of Sports Medicine, Peking University Third Hospital, Beijing, China

³ Department of Neurology, Peking Union Medical College Hospital, Beijing, China
zhuyc@pumch.cn

⁴ Alder Hey Children's NHS Foundation Trust, L14 5AB, Liverpool, UK

Abstract. Accurate 3D bone defect completion is critical for personalized implant design and surgical planning, where both boundary alignment and surface quality are essential. Existing voxel-, point-, or implicit-based generative models often struggle to simultaneously preserve topology, surface fidelity, and conditional consistency near defect regions. We propose FDMMBone, a flow-based conditional mesh completion framework operating in a deformable tetrahedral mesh (DMTet) parameter space. By representing shapes on a unified, geometry-aware tetrahedral grid, FDMMBone enables stable conditional generation of mesh parameters via probability flow ODE integration. A subsequent subdivision-based deformation refinement further enhances surface smoothness and boundary continuity. Experiments on multiple public datasets demonstrate consistent improvements over state-of-the-art methods in shape accuracy, boundary fidelity, and generalization, highlighting the robustness and clinical potential of the proposed approach.

Keywords: Flow-based model · Deformable tetrahedral mesh · 3D shape generation · Defective bone repair.

1 Introduction

3D bone shape completion is a core problem in medical image analysis and surgical planning, where inaccurate reconstruction can lead to poor implant fitting and increased surgical risks [10]. Clinically and geometrically, reliable completion requires recovering anatomically plausible global structure while ensuring precise boundary alignment and smooth surfaces, which remains challenging under complex defects and incomplete observations [30].

Early bone completion methods mainly relied on traditional techniques [14], which often impose strong assumptions on defect types and shapes, thus generalizing poorly to complex or large-scale defects. Recent deep learning ap-

proaches have explored voxel [19,6], point cloud [5], implicit SDF [21], symmetry-aware [25], and data-augmented generative frameworks [26] for 3D bone and cranial defect completion. However, voxel-based methods incur high computational costs at increased resolutions, point clouds lack explicit topology, and implicit SDF representations require costly continuous regression and sampling. In addition, symmetry-based and heavily augmented methods can improve robustness, but they may still be constrained by their reliance on anatomical priors and the difficulty of preserving topological consistency. More compact representations, such as triplanes or Gaussian-based models, improve efficiency but still struggle to maintain accurate surface geometry and topological consistency.

Flow matching learns a continuous velocity field and generates samples by probability flow ODE integration, offering better stability under strong geometric conditioning than diffusion-based noise prediction. To adapt it to mesh generation, we use a deformable tetrahedral parameterization based on Deep Marching Tetrahedra (DMTet) [20], which embeds surface geometry into a fixed-dimensional parameter space. This enables efficient coarse mesh generation in the first stage, followed by subdivision and geometric refinement to recover fine surface details. Our contributions are summarized below:

1. We propose a conditional flow matching formulation that transports noise to anatomically valid implant parameters in a deformable tetrahedral space.
2. We introduce a geometry-constrained projection-correction within probability flow ODEs to enforce anatomical boundary consistency during sampling.
3. We design a coarse-to-fine mesh completion framework that separates global shape generation from local refinement for accurate implant reconstruction.

2 Method

Pipeline FDMMBone adopts a coarse-to-fine framework for triangular mesh completion from defective bone inputs (Fig. 1). Both the defective mesh C and the target implant mesh Y are first mapped into a DMTet parameter space, where each vertex on a fixed tetrahedral grid is represented by a displacement vector and an SDF value. In the first stage, a conditional flow model is trained and probability flow ODE integration is used to generate shape parameters, from which a coarse mesh is extracted via marching tetrahedra. In the second stage, the coarse mesh is subdivided and refined through deformation-based adjustment, yielding a smooth and geometrically consistent final mesh.

2.1 Tetrahedral mesh parameterization

We adopt DMTet [20] to parameterize 3D bone shapes by discretizing the space into tetrahedral elements. Each vertex carries a position offset and a SDF, allowing explicit representation of bone surfaces while maintaining topological flexibility. To obtain a regular parameter space, we initialize the tetrahedral mesh using a Body-Centered Cubic (BCC) tiling scheme [8,15], which provides good spatial uniformity, near-regularity, and consistent interior vertex valence.




Fig. 1: Overview of the proposed framework. Coarse completion is generated by conditional flow matching in the DMTet parameter space and refined by a deformation network, with projection-correction applied during ODE integration.

In all experiments, we use a fixed DMTet grid with resolution 128^3 , resulting in $N = 277410$ tetrahedral vertices. The DMTet representation is obtained by deterministic fitting rather than by an additional preprocessing network. For each training mesh, we optimize vertex offsets and SDF values using surface Chamfer distance, SDF regularization, and Laplacian smoothness, producing aligned DMTet parameter pairs for the defective bone and the target implant. Let the initial tetrahedral grid have vertices $\{\mathbf{v}_i^0\}_{i=1}^N$. In DMTet, each vertex is associated with a displacement and an SDF value: $\mathbf{a}_i = [\Delta\mathbf{v}_i; s_i] \in \mathbb{R}^4$, where $\Delta\mathbf{v}_i \in \mathbb{R}^3$ and $s_i \in \mathbb{R}$. Stacking all vertex attributes yields the DMTet parameter representation: $x = [\mathbf{a}_1; \mathbf{a}_2; \dots; \mathbf{a}_N] = [\Delta\mathbf{V}; \mathbf{s}] \in \mathbb{R}^{N \times 4}$, where $\Delta\mathbf{V} \in \mathbb{R}^{N \times 3}$ collects $\{\Delta\mathbf{v}_i\}_{i=1}^N$ and $\mathbf{s} \in \mathbb{R}^{N \times 1}$ collects $\{s_i\}_{i=1}^N$. Similarly, the defective bone is encoded on the same grid as $c = [\Delta\mathbf{V}^c; \mathbf{s}^c] \in \mathbb{R}^{N \times 4}$.

2.2 Flow matching model in the DMTet parameter space

Flow modeling In bone defect repair, the defective bone C is used as the conditional input, and the objective is to generate an implant consistent with C . Both the defective bone C and the target implant mesh Y are represented in the DMTet parameter space. We denote the clean implant parameters as $x_1 \in \mathbb{R}^{N \times 4}$ and the conditional parameters of the defective bone as $c \in \mathbb{R}^{N \times 4}$. The goal of the first stage is to learn the conditional distribution $p(x_1 | c)$.

Since DMTet parameters are not directly available, the implant mesh Y and the defective bone C are first fitted to the DMTet representation under full 3D supervision to obtain aligned parameter pairs (x_1, c) using a geometric consistency loss [1]. Based on these aligned pairs, we train a conditional flow matching model in the DMTet parameter space. Following the standard flow matching

formulation [13], we sample a noise endpoint $x_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and define a linear interpolation path between the noise and the clean implant parameters as: $x(t) = (1-t)x_0 + tx_1$, $t \sim \text{Unif}(0, 1)$. The target velocity induced by this path is given by $v(x(t), t) = x_1 - x_0$, which is constant for each sampled pair (x_0, x_1) . A conditional network $\hat{v}_\theta(x(t), t, c)$ is trained to regress this velocity using a weighted ℓ_2 loss [29]. The conditional velocity field network follows a U-shaped encoder-decoder architecture. Each encoder block consists of two residual modules followed by a downsampling operation, while each decoder block contains two residual modules and an upsampling layer. Within each residual module (Fig. 2), layer normalization (LN), self-attention (SA), and cross-attention (CA) are employed, where SA captures global latent dependencies and CA injects conditional information from the defective bone.

During inference, sampling starts from the noisy latent state $x_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and progressively evolves towards the clean implant parameters x_1 by integrating the probability flow ODE over the normalized time interval $t \in [0, 1]$:

$$\frac{dx(t)}{dt} = \hat{v}_\theta(x(t), t, c), \quad t \in [0, 1]. \quad (1)$$

The coarse implant mesh is then extracted via the DMTet isosurface operation.




Fig. 2: The residual module of velocity field network.

Projection correction strategy During probability flow ODE integration, we introduce a lightweight projection-correction strategy to reduce local instabilities stemming from discretization errors and model inaccuracies while improving consistency between the generated shape and the defective condition. After obtaining an intermediate state $\tilde{x}_{i-1}$, we apply a proximal projection to ensure that the corrected state x_{i-1} remains close to $\tilde{x}_{i-1}$ and lies within a geometrically feasible, condition-consistent set. The projection is formulated as a regularized optimization problem:

$$x_{i-1} = \arg \min_x \frac{1}{2\lambda_i} \|x - \tilde{x}_{i-1}\|_2^2 + \mathcal{L}_{\text{cond}}(x | c) + \mathcal{L}_{\text{lap}}(x), \quad (2)$$

Here, $\mathcal{L}_{\text{cond}}(x | c)$ [3] is computed in a narrow band around the defect boundary and penalizes inconsistency between the generated implant and the observed defective bone. $\mathcal{L}_{\text{lap}}(x)$ [24] is a local Laplacian regularization term on the DMTet vertex offsets, which suppresses high-frequency artifacts. In all experiments, we set $\lambda_i = 0.1$, use $K = 2$ correction steps, and set the gradient step size to $\eta = 1 \times 10^{-3}$. We approximate the solution using a few gradient descent steps:

$$x^{(k+1)} = x^{(k)} - \eta \nabla_x E_i(x^{(k)}), \quad x^{(0)} = \tilde{x}_{i-1}, \quad k = 0, \dots, K-1. \quad (3)$$

The corrected state $x_{i-1} = x^{(K)}$ is obtained after K steps (with $K = 2$ in our experiments). This projection-correction reduces instabilities along the sampling path with minimal additional cost, resulting in smoother and more stable DM Tet parameters, better aligned with defect-conditioned constraints for reliable isosurface extraction.

2.3 Geometric refinement network

Based on the DM Tet parameters generated by the first-stage flow model, we extract a coarse implant mesh $M_c = (\mathbf{V}_c, \mathbf{F}_c)$ via an isosurface operation. Here, $\mathbf{V}_c$ is implicitly determined by the DM Tet vertex offsets ΔV_{stage1} predicted in the parameter space. To improve local smoothness and recover finer geometric details without increasing the tetrahedral resolution, we introduce a lightweight second-stage geometric refinement module.

Specifically, we first apply face subdivision to M_c to obtain a higher-resolution mesh $M_s = (\mathbf{V}_s, \mathbf{F}_s)$. A mesh deformation network then predicts per-vertex refinement displacements $\Delta \mathbf{V}_{\text{refine}}$, producing the refined mesh $M_r = (\mathbf{V}_r, \mathbf{F}_s)$ with: $\mathbf{V}_r = \mathbf{V}_s + \Delta \mathbf{V}_{\text{refine}}$. Conditional information from the defective bone is injected into the refinement network to preserve geometric consistency near defect boundaries. The refinement network is trained under supervised geometric constraints. We enforce global shape consistency using the chamfer distance $\mathcal{L}_{\text{surf}}$ [12], introduce a normal consistency loss $\mathcal{L}_n$ [17] to encourage alignment of surface orientations, and apply a vertex-based Laplacian regularization $\mathcal{L}_{\text{lap}}$ [23,24] to suppress high-frequency noise and improve mesh quality. The overall refinement loss is defined as: $\mathcal{L}_{\text{ref}} = \mathcal{L}_{\text{surf}} + w_n \mathcal{L}_n + w_l \mathcal{L}_{\text{lap}}$. This lightweight refinement strategy significantly improves surface smoothness and local details while keeping the reasoning cost low. Near defect boundaries, it produces more continuous and better-fitting geometric reconstructions.

3 Experiments

Implementation details We initialize a unified tetrahedral mesh $\mathcal{T}$ in the cubic domain $[-1, 1]^3$ using a dense BCC arrangement [8]. The allowable vertex deformation range is set to $2\times$ that of [20] to better capture high-frequency geometric details. Vertices not contributing to surface generation are assigned fixed signed distance values of ± 1 based on their inside-outside status. The flow model uses a 3D U-Net with a base width of 128, with defect conditions injected via conditional modulation. It is trained using AdamW with a learning rate of 1×10^{-4} for 6×10^4 iterations and a batch size of 12. During inference, samples are generated by solving the probability flow ODE from x_0 using a fixed-step solver with 200 steps, followed by the marching tetrahedra isosurface extraction. We use explicit Euler with step size $\Delta t = 1/200$. After each solver step, we apply the proximal projection to obtain the corrected state. The refinement network adopts a U-Net architecture composed of three residual blocks and is trained for 200 epochs with 5×10^{-4} learning rate. The full model is trained for 33 hours

Table 1: Evaluation of DSC(%), bDSC(%), HD95(mm) and MMD on MedShapeNet, Mandibular and SkullBreak datasets. The MMD values are multiplied by 10^2 . Best results in bold, second-best results in underlined.

Method	MedShapeNet				Mandibular				SkullBreak			
	DSC↑	bDSC↑	HD95↓	MMD↓	DSC↑	bDSC↑	HD95↓	MMD↓	DSC↑	bDSC↑	HD95↓	MMD↓
SDFusion	83.3	83.4	3.19	0.344	83.4	83.9	3.17	0.349	83.7	83.0	3.15	0.353
Diffusion-SDF	85.7	83.0	3.01	0.309	84.2	82.1	2.91	0.307	85.4	83.8	2.86	0.305
DiffComplete	86.9	87.2	1.94	0.220	87.7	86.9	1.87	0.216	88.1	88.2	1.77	0.201
Pcdiff	88.6	91.5	1.98	0.211	88.5	91.7	1.90	0.215	88.2	91.2	2.11	0.213
OctFusion	89.1	89.4	1.89	0.213	89.5	88.9	1.85	0.209	89.7	89.6	1.82	0.208
MeshGPT	91.3	91.1	<u>1.70</u>	0.215	90.6	91.0	<u>1.79</u>	<u>0.192</u>	91.1	90.5	<u>1.78</u>	<u>0.181</u>
LT3SD	<u>93.1</u>	<u>93.2</u>	1.71	<u>0.185</u>	<u>92.7</u>	<u>92.8</u>	1.78	0.193	<u>92.4</u>	<u>93.0</u>	1.80	0.182
Ours	94.4	94.7	1.51	0.113	94.3	94.8	1.44	0.119	94.6	94.5	1.37	0.111

on two NVIDIA A800 GPUs. During inference, generating one implant takes about 2.9 seconds, where the 200-step ODE solver takes about 2.1 seconds and the projection-correction plus refinement stages take about 0.8 seconds.

Dataset and evaluation metrics We conduct experiments on four public datasets. MedShapeNet [11] includes 3870 vertebral samples, and the Mandibular dataset [27] contains 145 real patient cases with mandibular bone defects; both datasets are split into training and testing sets with a 4:1 ratio. SkullBreak [7] follows the official split with 570 training cases and 100 testing cases. The MUG500 dataset [9] is used only for cross-dataset evaluation, where the 29 defective cases are adopted to assess the generalization capability of our method. Following prior works [5,16], we report four standard evaluation metrics, including DSC, bDSC, HD95, and MMD.

3.1 Bone Completion Results and Discussion

Alternative Methods. We compared FDMMBone with seven SOTA automatic shape reconstruction and shape completion methods, including SDFusion[2], Diffusion-SDF[3], DiffComplete[4], Pcdiff[5], Octfusion [28], MeshGPT[22] and LT3SD[18]. For fair comparison, all competing methods are retrained using the same train/test splits, preprocessing pipeline, and evaluation metrics.

Comparison results Fig. 3 shows a comparison of reconstructed implants using 3D signed distance error colormaps, with errors normalized to $[-3, 3]$ and generated by MeshMetric3D tool: <https://www.nitrc.org/projects/meshmetric3d/>.

SDFusion [2] shows deviations near defect boundaries and high-curvature regions, likely due to the smooth nature of implicit fields. Diffusion-SDF [3] exhibits mild averaging at sharp boundaries and thin structures, with errors concentrated around complex defect edges. DiffComplete [4] provides globally stable reconstructions but still shows scattered local deviations. Pcdiff [5] produces discontinuous color bands near boundaries, likely due to the lack of explicit connectivity in point clouds. OctFusion [28] is stable overall but shows local variations in high-curvature regions, indicating limited cross-level detail




Fig. 3: Qualitative comparison among the complete methods implemented on three datasets. These colormaps depict the signed distance between the reconstructed implants and the ground truth bone meshes.

propagation. MeshGPT [22] has slight color shifts near defect boundaries, possibly related to boundary alignment limits in serialized mesh representations. LT3SD [18] achieves coherent global reconstructions but retains higher deviations around fine boundaries. In contrast, FDMMBone shows closer alignment with the ground truth, fewer deviations, and smoother transitions, benefiting from flow-based modeling on a fixed tetrahedral grid and a refinement stage that improves local surface continuity.

As shown in Table 1, our method consistently achieves the best performance across all datasets, especially in DSC and HD95. Higher DSC indicates more accurate global shape recovery, while lower HD95 reflects improved boundary alignment. Compared with diffusion- or point-based methods, our approach better balances global anatomical consistency and boundary-level geometric fidelity, which is essential for clinically reliable bone completion.

Generalization ability To evaluate cross-dataset generalization, we train the model on SkullBreak and directly test it on the unseen MUG500 dataset.




Fig. 4: Generalization experiment involving the proposed method.

Quantitatively, FDMMBone achieves a DSC of 91.8% and bDSC 91.7%. Moreover, the method maintains strong geometric fidelity, with an HD95 of 1.72 and a low MMD of 0.122. As shown in Fig. 4, our method produces high-fidelity completions with consistent structure and smooth surfaces despite large distribution shifts. Boundary transitions remain natural, with no obvious performance degradation compared to in-domain results in Table 1. Notably, this is achieved without any domain adaptation or fine-tuning, demonstrating strong robustness and transferable anatomical priors. Since FDMMBone outputs explicit triangular meshes, the generated implants can be directly exported as STL/OBJ files and imported into standard medical CAD software, such as 3D Slicer or 3-matic, for non-manifold checking, Boolean validation, and manufacturing preparation.

Table 2: Ablation study of FDMMBone on SkullBreak dataset.

Variant	DSC↑	bDSC↑	HD95↓	MMD↓
(A1) w/o Stage-2 refinement (only Stage-1 flow model)	93.4	93.5	1.51	0.121
(A2) w/o Subdivision (refiner deforms M_c directly)	93.6	93.7	1.49	0.117
(A3) Loss of Stage-2: only $\mathcal{L}_{\text{surf}}$	93.7	93.9	1.46	0.118
(A4) Loss of Stage-2: $\mathcal{L}_{\text{surf}} + \mathcal{L}_n$	93.8	94.0	1.40	0.115
(A5) Loss of Stage-2: $\mathcal{L}_{\text{surf}} + \mathcal{L}_{\text{lap}}$	94.0	94.1	1.46	0.114
(A6) w/o Conditional injection (no c features)	91.3	91.2	1.67	0.172
(A7) $K = 0$ & $(w_n, w_l) = (0.7, 0.3)$	93.2	93.1	1.70	0.139
$K = 2$ & $(w_n, w_l) = (0.5, 0.5)$	94.5	94.4	1.49	0.113
$K = 2$ & $(w_n, w_l) = (0.3, 0.7)$	94.4	94.2	1.53	0.119
Ours (w/ $K = 2$ & $(w_n, w_l) = (0.7, 0.3)$)	94.6	94.5	1.37	0.111

3.2 Ablation study

As shown in Table 2, removing conditional injection (A6) results in the most significant performance drop, especially in bDSC and HD95, indicating the importance of defect-conditioned guidance near boundaries. Removing the Stage-2 refinement module (A1) also degrades performance, confirming the benefit of the coarse-to-fine design. Within the refinement stage, disabling mesh subdivision (A2) limits the representation of local details, while removing conditional injection (A6) mainly degrades boundary-related metrics such as bDSC and HD95, indicating the necessity of explicit conditional guidance near defect boundaries. The normal consistency and Laplacian terms (A3/A4/A5) primarily contribute to suppressing local noise and stabilizing surface quality, leading to consistent

improvements in HD95 and MMD. For projection-correction, setting $K = 0$ disables the correction step, while $K = 2$ improves all metrics, showing that the correction stabilizes ODE sampling. For the refinement loss, $(w_n, w_l) = (0.7, 0.3)$ achieves the best balance between boundary accuracy and surface smoothness.

4 Conclusion

We propose FDMMBone for accurate medical bone completion in triangle-mesh form. Our method performs conditional flow matching in a DM Tet space with a lightweight projection-correction step for stable, condition-consistent sampling, followed by subdivision-based deformation refinement. Experiments show consistent improvements across metrics and strong generalization.

Acknowledgments. This research was partially supported by the Beijing Municipal Science and Technology Commission and Zhongguancun Science Park Management Committee (no. Z221100002722020), the National Nature Science Foundation of China (no. 62072045), and the National Nature Science Foundation of Beijing (no. 7242167).

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

References

1. Cao, Y., Chen, B., Li, X., Liang, Y., Sha, Z., Shi, Z., Song, Z., Wan, M.: Force matching with relativistic constraints: A physics-inspired approach to stable and efficient generative modeling. In: Proceedings of the 34th ACM International Conference on Information and Knowledge Management. pp. 179–188 (2025)
2. Cheng, Y.C., Lee, H.Y., Tulyakov, S., Schwing, A.G., Gui, L.Y.: Sdfusion: Multi-modal 3D shape completion, reconstruction, and generation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4456–4465 (2023)
3. Chou, G., Bahat, Y., Heide, F.: Diffusion-SDF: Conditional generative modeling of signed distance functions. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 2262–2272 (2023)
4. Chu, R., Xie, E., Mo, S., Li, Z., Nießner, M., Fu, C.W., Jia, J.: Diffcomplete: Diffusion-based generative 3D shape completion. *Advances in neural information processing systems* **36**, 75951–75966 (2023)
5. Friedrich, P., Wolleb, J., Bieder, F., Thieringer, F.M., Cattin, P.C.: Point cloud diffusion models for automatic implant generation. In: Medical Image Computing and Computer Assisted Intervention. pp. 112–122 (2023)
6. Huang, Y., Zou, S., Liu, X., Xu, K.: Part-aware shape generation with latent 3d diffusion of neural voxel fields. *IEEE Transactions on Visualization and Computer Graphics* **31**(10), 8057–8069 (2025)
7. Kodym, O., Li, J., Pepe, A., Gsaxner, C., Chilamkurthy, S., Egger, J., Španěl, M.: Skullbreak/skullfix–dataset for automatic cranial implant design and a benchmark for volumetric shape learning tasks. *Data in Brief* **35**, 106902 (2021)

8. Labelle, F., Shewchuk, J.R.: Isosurface stuffing: fast tetrahedral meshes with good dihedral angles. In: ACM SIGGRAPH 2007 papers, pp. 57–es (2007)
9. Li, J., Krall, M., Trummer, F., Memon, A.R., Pepe, A., Gsaxner, C., et al.: Mug500+: Database of 500 high-resolution healthy human skulls and 29 craniotomy skulls and implants. Data in Brief **39**, 107524 (2021)
10. Li, J., Von Campe, G., Pepe, A., Gsaxner, C., Wang, E., Chen, X., et al.: Automatic skull defect restoration and cranial implant generation for cranioplasty. Medical Image Analysis **73**, 102171 (2021)
11. Li, J., Zhou, Z., Yang, J., Pepe, A., Gsaxner, C., Luijten, G., et al.: MedShapeNet—a large-scale dataset of 3D medical shapes for computer vision. Biomedical Engineering/Biomedizinische Technik **70**(1), 71–90 (2025)
12. Lin, F., Yue, Y., Zhang, Z., Hou, S., Yamada, K., Kolachalama, V., Saligrama, V.: Infocd: A contrastive chamfer distance loss for point cloud completion. Advances in Neural Information Processing Systems **36**, 76960–76973 (2023)
13. Lipman, Y., Chen, R.T.Q., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: International Conference on Learning Representations (2023)
14. Liu, J., Savva, M., Mahdavi-Amiri, A.: Survey on modeling of human-made articulated objects. In: Computer Graphics Forum. p. e70092 (2025)
15. Liu, Z., Feng, Y., Black, M.J., Nowrouzezahrai, D., Paull, L., Liu, W.: MeshDiffusion: Score-based generative 3D mesh modeling. arXiv preprint arXiv:2303.08133 (2023)
16. Liu, Z., Ru, X., Wang, X., Wu, Z., Zhu, Y.C., Zhang, C., Frangi, A.F.: 3d skull completion via two-stage conditional diffusion-based signed distance fields. In: IEEE International Conference on Bioinformatics and Biomedicine. pp. 2204–2209 (2024)
17. Men, Y., Lei, B., Yao, Y., Cui, M., Lian, Z., Xie, X.: En3d: An enhanced generative model for sculpting 3D humans from 2D synthetic data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9981–9991 (2024)
18. Meng, Q., Li, L., Nießner, M., Dai, A.: Lt3sd: Latent trees for 3D scene diffusion. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 650–660 (2025)
19. Romanelis, I., Fotis, V., Kalogeras, A., Alexakos, C., Munteanu, A., Moustakas, K.: Efficient and scalable point cloud generation with sparse point-voxel diffusion models. IEEE Transactions on Neural Networks and Learning Systems **99**, 1–12 (2025)
20. Shen, T., Gao, J., Yin, K., Liu, M.Y., Fidler, S.: Deep marching tetrahedra: a hybrid representation for high-resolution 3D shape synthesis. Advances in Neural Information Processing Systems **34**, 6087–6101 (2021)
21. Shim, J., Kang, C., Joo, K.: Diffusion-based signed distance fields for 3D shape generation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 20887–20897 (2023)
22. Siddiqui, Y., Alliegro, A., Artemov, A., Tommasi, T., Sirigatti, D., Rosov, V., Dai, A., Nießner, M.: Meshgpt: Generating triangle meshes with decoder-only transformers. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 19615–19625 (2024)
23. Wang, J., Wu, S., Liu, A., Meng, Z.: Talent3D: Optimizing geometry and enhancing appearance for high-quality Text-to-3D shape generation. Neurocomputing **645**, 130475 (2025)

24. Wang, N., Zhang, Y., Li, Z., Fu, Y., Liu, W., Jiang, Y.G.: Pixel2mesh: Generating 3D mesh models from single RGB images. In: Proceedings of the European conference on computer vision (ECCV). pp. 52–67 (2018)
25. Wodzinski, M., Daniol, M., Hemmerling, D.: Automatic skull reconstruction by deep learnable symmetry enforcement. *Computer Methods and Programs in Biomedicine* **263**, 108670 (2025)
26. Wodzinski, M., Kwarciak, K., Daniol, M., Hemmerling, D.: Improving deep learning-based automatic cranial defect reconstruction by heavy data augmentation: From image registration to latent diffusion models. *Computers in biology and medicine* **182**, 109129 (2024)
27. Wu, J., Jiang, L., Shao, L., Wang, W., Xu, X., Zhou, Y., et al.: A mandibular defect dataset for autonomous reconstruction planning in oral and maxillofacial surgery. *Scientific Data* **12**(1), 1763 (2025)
28. Xiong, B., Wei, S.T., Zheng, X.Y., Cao, Y.P., Lian, Z., Wang, P.S.: Octfusion: Octree-based diffusion models for 3d shape generation. In: *Computer Graphics Forum*. vol. 44, p. e70198 (2025)
29. Yazdani, M., Medghalchi, Y., Ashrafian, P., Hacıhaliloglu, I., Shahriari, D.: Flow matching for medical image synthesis: Bridging the gap between speed and quality. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 216–226 (2025)
30. Zhou, J., See, C.W., Sreenivasamurthy, S., Zhu, D.: Customized additive manufacturing in bone scaffolds—the gateway to precise bone defect treatment. *Research* **6**, 0239 (2023)