

Shape-guided Gaussian Splatting for Sparse-View X-ray 3D Reconstruction

Pranav Poudel^{1,3*}, Florence Dell’Aniello Picard^{1,3}, Nairouz Shehata^{1,3},
Frédéric Lavoie², and Herve Lombaert^{1,3}

¹ Polytechnique Montréal, Canada

² CHUM - University of Montreal Hospital, Canada

³ Mila - Quebec AI Institute, Canada

Abstract. Sparse-view X-ray 3D reconstruction is essential for reducing radiation exposure, but recovering a density field from a handful of X-ray projections is severely ill-posed. Recently, 3D Gaussian Splatting has achieved state-of-the-art performance in sparse-view reconstruction by representing the volume using explicit, optimized primitives, but it requires dozens of projected views. With fewer views, reconstruction quality degrades severely since the explicit primitives are optimized freely without any anatomical information. Anatomical structures, in contrast, share similar geometry and density across a population. Their variations are bounded within a limited range that statistical shape models can capture. This paper proposes a shape-guided Gaussian splatting framework for sparse-view X-ray 3D reconstructions. Our contribution lies in driving Gaussian positions toward anatomically valid configurations, alongside atlas-based density regularization. Our method ensures anatomically consistent reconstruction and improves PSNR by 2.83 dB over a state-of-the-art Gaussian splatting baseline with as few as 5 views. Code Available: <https://github.com/polyshape-lab/ShapeGuidedGaussian>

Keywords: Gaussian splatting · 3D reconstruction · Shape analysis

1 Introduction

3D reconstructions of human anatomy are relevant to clinical applications, particularly in orthopedics, such as preoperative planning, surgical navigation, and patient-specific instrumentation [12,23]. X-ray imaging is widely used in routine clinical practice. However, recovering 3D geometry from a few 2D X-ray projections is an ill-posed problem. CT imaging can reconstruct a high-resolution 3D volume by acquiring hundreds of projections and inverting the Radon transform [18], but at a much higher radiation dose than standard X-ray imaging [10]. Reducing the number of projections lowers this dose and eases acquisition in clinical settings such as the operating room. Prior work on recovering 3D anatomy from X-ray projections falls into two families: surface reconstruction [4,3,17,19]

* Correspondence to: Pranav Poudel <pranav.poudel@etud.polymtl.ca>

and volumetric reconstruction [30,31,9,14,1]. Surface reconstruction methods recover the 3D boundary using a statistical shape model, but rely on features such as accurate contours and landmarks of its boundary. Volumetric reconstructions recover a density field, i.e. the X-ray attenuation at each point in the object, directly from X-ray projections, but require dozens of projections. We focus on volumetric reconstructions, often referred to as tomographic reconstructions.

Traditional tomographic reconstruction methods [14,1] produce high-quality reconstructions with hundreds of projections, but quality degrades severely under sparse views, leading to streak artifacts and loss of structural detail. More recently, NeRF-based methods [31,9], and 3D Gaussian splatting methods [8,30] have drawn interest for volumetric reconstruction. NeRF-based methods model the density field with multilayer perceptrons, whereas Gaussian splatting uses explicit Gaussian primitives; both are trained with photometric loss. Among these, R^2 -Gaussian [30] achieves state-of-the-art performance in sparse-view settings (25–75 views) by addressing the integration bias of 3DGS [20]. However, both families predict the density field solely from projections, with no anatomical prior during optimization. Under extremely sparse views, reconstructions become severely ill-posed, with many density fields that can project to the same view. These methods then overfit the available views, producing ill-formed boundaries.

Human anatomy shares similar geometry across individuals, with variation bounded within a limited range [11]. Furthermore, organs have extensive homogeneous regions within but change sharply at their boundaries. Statistical shape models capture such regularities and are widely used in a variety of medical imaging applications [16]. Shape models based on Gaussian Splatting have been used in computer vision, notably for faces and human avatars [32,22]. However, such methods require large numbers of images from video clips, which is often infeasible in a medical setting. The bounded variation captured by a shape model can serve as a strong prior that can meet the requirements of many 2D projections in medical settings. Integrating shape models into Gaussian splatting remains underexplored for medical imaging. Shape-based priors have been used to constrain 3D reconstruction from a few projections, through learned structure priors [7] and atlas-based 2D/3D registration [27]. To our knowledge, no existing work binds Gaussian primitives to a shape model for X-ray 3D reconstruction.

To overcome the lack of anatomical constraints in existing sparse-view 3D reconstruction methods, a shape-guided Gaussian splatting framework is proposed that uses population shape and density information to guide reconstruction. Our guidance is applied at two stages. *During initialization*, Gaussian positions and orientations are seeded from the shape template, and Gaussian densities are seeded from the population-average density from the intensity image. *During optimization*, Gaussian positions are collectively deformed within an anatomically plausible region via a shape prior, while a density prior constrains the reconstructed volume to statistical variations in density. Our method substantially outperforms the state-of-the-art baseline when reconstructing femoral shapes under extremely sparse views. Our contributions are as follows:




Fig. 1. Overview of our framework. Shape Gaussians (Γ_s) are initialized from a template and population mean density, seeding reference position $\mathbf{p}^{\text{ref}}$, scaling S , rotation R and density ρ . Positions deform collectively via a displacement field $\mathbf{u}(\alpha)$ generated by the shape coefficients α , while scaling, rotation, and density are updated independently. The full Gaussian primitives set ($G = \Gamma_s \cup \Gamma_z$) are projected by the X-ray rasterizer, and all parameters are optimized with photometric ($\mathcal{L}_1$, $\mathcal{L}_{\text{ssim}}$) and regularization ($\mathcal{L}_{\text{tv}}$, $\mathcal{L}_{\text{density}}$, $\mathcal{L}_{\text{shape}}$) losses.

- We propose a shape-guided Gaussian splatting framework that binds the positions of Gaussian primitives to a shape model, constraining the geometry to anatomically valid configurations.
- We incorporate guidance in two stages: Initialization and Optimization. We seed orientation and appearance of primitives from the population mean and regularize the reconstruction towards population statistics.
- We conduct experiments on femoral data from NMDID [13] and demonstrate consistent improvements over the state-of-the-art baseline under extreme sparsity as low as 5-view reconstruction.

2 Method

The proposed method guides Gaussian splatting using population priors to ensure that reconstruction remains anatomically plausible. As shown in Fig. 1, the guidance occurs in two stages: Initialization and Optimization. In this section, the object representation is introduced first (Sec. 2.1), followed by the shape model that guides the Gaussian primitives (Sec. 2.2). The Parameterization and Initialization of Gaussian primitives using the shape model are elaborated (Sec. 2.3), followed by their optimization (Sec. 2.4).

2.1 Object Representations

The 3D volume can be represented as a collection of N radiative Gaussian primitives [30] $\{G_i\}_{i=1}^N$ each defined as,

$$G_i(\mathbf{x}) = \rho_i \cdot \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{p}_i)^\top \Sigma_i^{-1}(\mathbf{x} - \mathbf{p}_i)\right), \quad (1)$$

with position $\mathbf{p}_i \in \mathbb{R}^3$, covariance Σ_i , and central density $\rho_i \in \mathbb{R}_{\geq 0}$. For optimization purposes [20], the covariance matrix is decomposed as $\Sigma_i = R_i S_i S_i^\top R_i^\top$, where R_i is the rotation and S_i is the scaling matrix. This collection of primitives is partitioned into two sets: a shape-guided set Γ_s , representing the anatomy of interest that will be guided by the shape model, and a free set Γ_z representing the rest of the volume. Thus, primitives in Γ_z have position, covariance, and density all as free learnable parameters, whereas Gaussians in Γ_s have their position driven by a shape model. Since Gaussian primitives combine additively, the density field at ray position $\mathbf{x}$ is defined as $\mu(\mathbf{x}) = \sum_{i=1}^N G_i(\mathbf{x})$. A differentiable rasterizer renders a projection I_r as an analytic line integral of the attenuation field along each ray, and the voxelizer queries the field on a grid to produce a density volume for 3D reconstruction [30].

2.2 Shape Model

The 3D geometry of an object is represented in a low-dimensional shape space learned from a population of P segmentations. Each subject in the population is registered to a common template using diffeomorphic registration [2], yielding a per-subject velocity field $\mathbf{v}$. A PCA over these velocity fields gives a mean velocity field $\bar{\mathbf{v}}$ and orthonormal modes $\{\phi_m\}$ with variance $\{\lambda_m\}$, so that a velocity field is generated from M shape coefficients, $\boldsymbol{\alpha} \in \mathbb{R}^M$,

$$\mathbf{v}(\boldsymbol{\alpha}) = \bar{\mathbf{v}} + \sum_{m=1}^M \alpha_m \sigma_m \phi_m, \quad \sigma_m = \sqrt{\lambda_m} \quad (2)$$

where standardizing by σ_m induces the prior $\boldsymbol{\alpha} \sim \mathcal{N}(\mathbf{0}, I)$. The velocity field $\mathbf{v}(\boldsymbol{\alpha})$ is integrated to obtain a displacement field $\mathbf{u}(\boldsymbol{\alpha})$ which is used to deform shape Gaussian primitives (Eq. 3). From the same registered population, per-voxel mean $\bar{V}_\mu(\mathbf{x})$ and variance $\sigma_V^2(\mathbf{x})$ of the density are computed to form a density atlas. This atlas is used for density initialization and confidence weight in the density prior loss.

2.3 Parameterization and Initialization

The core of the proposed method is to bind the Gaussian primitives to the shape model. Each $G_j(\mathbf{x}) \in \Gamma_s$ sits at a fixed reference center $\mathbf{p}_j^{\text{ref}}$ on the template. Instead of moving freely, it is carried by a deformation induced by the shape coefficients $\boldsymbol{\alpha}$, where the resulting displacement field $\mathbf{u}(\boldsymbol{\alpha})$ is sampled at the reference position. So, the position of each primitive in Γ_s is given by,

$$\mathbf{p}_j = \mathbf{p}_j^{\text{ref}} + \mathbf{u}_j(\boldsymbol{\alpha}). \quad (3)$$

In contrast, the scaling S_j , rotation R_j and density ρ_j that encode orientation and appearance remain independent parameters optimized directly. This will result in Gaussian primitives being in anatomically valid positions, while their appearance and orientation remain adaptive to the available projections. Thus, the learnable set of shape Gaussian primitives is: $\Theta = \{\boldsymbol{\alpha}; \{S_j\}, \{R_j\}, \{\rho_j\}\}$.

Prior work initializes Gaussian primitives from a low-quality volume reconstructed using FDK [14]. Under extreme sparsity, this reconstruction becomes severely degraded, giving an unreliable starting point for optimization. Instead, the orientation of the shape Gaussian primitives is initialized from the template, and the density is initialized from the population mean density. Reference centers $\{\mathbf{p}_j^{\text{ref}}\}_{j=1}^{N_s}$, where $N_s = |\Gamma_s|$, are obtained by boundary-weighted sampling of the template, placing denser primitives near the boundary where sharp transitions require more primitives to represent. Scaling parameters are seeded from the nearest-neighbour spacing of reference positions [30,20]. However, unlike their isotropic identity-rotation initialization, the surface primitives are made anisotropic by thinning primitives along the normal by a factor $\tau \in (0, 1)$, and orienting them in order to align their thinned axis with the normal of the template surface. This reorientation causes the Gaussian primitives to be flat across thin surfaces [15]. Densities are initialized so that the accumulated field matches the per-voxel population mean density. Since primitives overlap and combine additively, a single primitive’s density cannot be read directly from the target. To account for this, all densities are set to one and voxelized to measure how much overlap occurs at each location using Eq. 4

$$n(\mathbf{x}) = \sum_{i \in \Gamma_s} \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{p}_i)^\top \Sigma_i^{-1}(\mathbf{x} - \mathbf{p}_i)\right). \quad (4)$$

Each density is then set to atlas mean density $\bar{V}_\mu$ at the Gaussian primitives reference location $\mathbf{p}_j^{\text{ref}}$, divided by the local overlap defined in Eq. (4) as:

$$\rho_j^{\text{init}} = \frac{\bar{V}_\mu(\mathbf{p}_j^{\text{ref}})}{\max(n(\mathbf{p}_j^{\text{ref}}), 1)}. \quad (5)$$

Hence, the overlapping Gaussians sum to the correct mean density.

2.4 Optimization Strategy

We optimize Gaussian primitives using gradient descent by minimizing the total objective L (Eq. 8). Besides combining photometric losses ($\mathcal{L}_1$ and D-SSIM $\mathcal{L}_{\text{ssim}}$ [28]) and a 3D total-variation prior ($\mathcal{L}_{\text{tv}}$ [24]), two regularization terms are introduced specific to the shape-guided setting: a shape prior on the shape coefficients and a density prior on the reconstructed volume. Shape modes are

standardized so $\boldsymbol{\alpha} \sim \mathcal{N}(\mathbf{0}, I)$, with plausible shapes lying near the origin of shape space. Hence, deviation is penalized as,

$$\mathcal{L}_{\text{shape}} = \|\boldsymbol{\alpha}\|_2^2, \quad (6)$$

which is the squared Mahalanobis distance of $\boldsymbol{\alpha}$ under shape prior. This keeps the deformed primitives within anatomically plausible shapes. A density prior pulls the reconstructed volume towards the population mean. But the population mean is only reliable where its value is consistent across subjects. On a randomly sampled sub-volume grid Ω , only shape Gaussian primitives are queried to obtain volume $\hat{V}$ and compute the density prior as:

$$w(\mathbf{x}) = \exp\left(-\frac{\sigma_V^2(\mathbf{x})}{s^2}\right), \quad \mathcal{L}_{\text{density}} = \frac{1}{|\Omega|} \sum_{\mathbf{x} \in \Omega} w(\mathbf{x}) (\hat{V}(\mathbf{x}) - \bar{V}_\mu(\mathbf{x}))^2, \quad (7)$$

where s controls the fall-off of the confidence weight w . Confidence is high where the population is consistent and low where it is highly variable. Hence, the final objective function becomes,

$$\mathcal{L} = \mathcal{L}_1 + \lambda_{\text{ssim}} \mathcal{L}_{\text{ssim}} + \lambda_{\text{tv}} \mathcal{L}_{\text{tv}} + \lambda_{\text{shape}} \mathcal{L}_{\text{shape}} + \lambda_{\text{density}} \mathcal{L}_{\text{density}}. \quad (8)$$

3 Experimental Setup

Datasets and Shape Model: We conduct experiments on synthetic datasets by leveraging 758 CT scans from NMDID [13], resampled into (1, 1, 1) mm spacing. TotalSegmentator [29] was used to generate femoral segmentations, and ANTs [26] was used for rigid alignment of all subjects to a common reference. Of the 758 subjects, 106 were held out: 100 to build a shape model and 6 for reconstruction. Since reconstruction is per-subject optimization, 6 subjects constitute 6 independent evaluations. The remainder was used to train VoxelMorph [5] for diffeomorphic registration. The trained model then co-registers the 100 subjects onto a common template, yielding a stationary velocity field for each subject. PCA over these fields was applied [25], retaining $M = 35$ components that capture 95% of variance. The voxelwise mean and variance over the same population were calculated for density initialization and the confidence weight in the density prior loss. We then use TIGRE [6] to synthesize X-ray projections over the 0–180° range, incorporating Compton scattering and electronic noise. To evaluate whether shape prior improves reconstruction quality under extreme sparsity, we set the number of views to 5 and 10, far sparser than the 25–75 views settings of prior work [30].

Baseline and metrics We compare our method against R^2 -Gaussian [30], a state-of-the-art 3D reconstruction method based on Gaussian splatting. We use PSNR and SSIM [28] to assess reconstruction quality, with PSNR calculated in 3D and SSIM averaged over 2D slices along 3 axes.

Implementation Details: Both our method and baseline are trained for 30K iterations with Adam [21] on a single A100 GPU. For photometric losses, we




Fig. 2. Qualitative comparison slices shown across two axes (top: coronal; bottom: sagittal). Compared to R^2 -Gaussian, our method recovers sharper cortical boundaries and finer internal structure, closer to the ground truth.


use the default hyperparameters of [30]. Proposed loss weights were selected as $\lambda_{\text{shape}} = 0.0001$, $\lambda_{\text{density}} = 0.5$, confidence fall-off $s = 0.05$, thinning factor $\tau = 0.5$ empirically on a single subject. Initial learning rates are $0.1(\alpha)$, $0.01(\rho)$, $0.005(\mathbf{S})$ and $0.01(\mathbf{R})$, each decaying exponentially to 10% of its initial value, except α , which decays to 1% by the midpoint of training. The shape coefficients are frozen during a short warm-up of 50 iterations to stabilize the appearance. We initialize $N_s = 35\text{K}$ shape Gaussians and $N_z = 15\text{K}$ free Gaussians (50K total matching baseline).

4 Results

Our experiments evaluate whether incorporating anatomical priors into Gaussian splatting improves reconstruction quality in extremely sparse-view settings. Table 1 shows the PSNR and SSIM for the baseline and our method across six subjects (**S1–S6**) at 5 and 10 views. Our results show that our method significantly outperforms the baseline, and the improvement is consistent across all subjects and both views. For example, our guidance improves the average 3D PSNR by +2.83 dB ($34.55 \rightarrow 37.38$) and 3D SSIM by +0.022 ($0.937 \rightarrow 0.959$) at 5 views, while at 10 views the gains are smaller: +1.46 dB ($38.89 \rightarrow 40.35$) and +0.012 ($0.959 \rightarrow 0.971$). The gap narrows as the number of views increases because the projections themselves provide more constraints. We further exam-

Table 1. Quantitative 3D results across six test subjects (S1–S6). Best in bold

# of Views	Metric	Method	S1	S2	S3	S4	S5	S6	Avg.
5	PSNR↑	R2GS	34.44	33.03	35.01	35.43	34.84	34.56	34.55
		Ours	37.71	35.69	37.93	37.48	37.41	38.06	37.38
	SSIM↑	R2GS	0.937	0.925	0.943	0.943	0.937	0.936	0.937
		Ours	0.961	0.948	0.965	0.960	0.957	0.961	0.959
10	PSNR↑	R2GS	39.08	38.05	39.51	39.58	38.68	38.44	38.89
		Ours	40.48	39.77	40.77	40.97	40.14	39.98	40.35
	SSIM↑	R2GS	0.961	0.955	0.965	0.964	0.956	0.954	0.959
		Ours	0.972	0.968	0.974	0.974	0.968	0.967	0.971

**Fig.3. Ablation Study on Initialization:** Free Gaussian (top) at 5 views vs. ground truth (bottom). Matched Projection at training view, yet the reconstructed volume does not match ground truth.**Table 2.** Ablation on 5-view reconstruction, showing the contribution of each proposed component.

Method	PSNR↑	SSIM↑
Density Init.	32.93	0.925
+ λ_{shape}	37.26	0.958
+ λ_{density}	37.38	0.959

ine whether incorporating the shape prior yields sharper, anatomically plausible boundaries than the baseline. Fig. 2 shows reconstructions of **S3** at 5 and 10 views along two different axes. At 5 views, we observe that the R^2 -Gaussian produces blurred boundaries, whereas our method recovers sharp boundaries and coherent structure. In the highlighted insets, we observe that the cortical boundary of the R^2 -Gaussian is smeared and spreads into the surrounding areas. In contrast, our method remains sharp and closely follows the ground-truth contour.

Ablation Study: We perform an ablation study to show the contribution of each proposed component. Table 2 and Fig. 3 show the contribution of each component on 5-view settings. Starting from density initialization alone, where all primitives are independent, adding the shape prior by binding primitives to the shape model yields a large improvement of +4.33 dB. Fig. 3 illustrates that with independent primitives, the rendered training-view angle projection closely matches the ground-truth projection, but the reconstructed volume does not. This is because only primitives along the observed rays are effectively optimized.

This results in a mottled, inconsistent interior despite having an initial boost. Adding the shape prior resolves this because all the primitives are deformed collectively. The density prior adds a further, smaller gain, refining the density.

5 Conclusion

We presented a shape-guided Gaussian splatting framework for sparse-view X-ray 3D reconstruction. This framework constrains the locations of Gaussian primitives to anatomically valid configurations using a shape model and leverages population statistics for regularization. Furthermore, our framework initializes the Gaussian parameters from population statistics, giving an anatomically valid starting point for optimization. Our method improves the PSNR by 2.83 dB and SSIM by 0.022 at only 5 views with respect to competitive state-of-the-art methods under extreme sparsity. Our method is, in theory, not limited to our tested femur, suggesting potential extensions to other anatomical structures.

Acknowledgments. This project is supported by the NSERC Alliance Advantage grant in partnership with Eiffel Medtech Inc. We acknowledge the Digital Research Alliance of Canada for providing computational resources, and the New Mexico Decedent Image Database (NMDID) for the CT imaging data.

Disclosure of Interests. This work has received research funding from Eiffel Medtech Inc. Dr. Lavoie is the founder of Eiffel Medtech Inc., and the other authors have no competing interests to declare.

References

1. Andersen, A.H., Kak, A.C.: Simultaneous algebraic reconstruction technique (sart): a superior implementation of the art algorithm. *Ultrasonic Imaging* **6** (1984)
2. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-euclidean framework for statistics on diffeomorphisms. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (2006)
3. Aubert, B., Vazquez, C., Cresson, T., Parent, S., de Guise, J.A.: Toward automated 3d spine reconstruction from biplanar radiographs using cnn for statistical spine model fitting. *IEEE Transactions on Medical Imaging* **38** (2019)
4. Baka, N., Kaptein, B.L., de Bruijne, M., van Walsum, T., Giphart, J., Niessen, W.J., Lelieveldt, B.P.: 2d–3d shape reconstruction of the distal femur from stereo x-ray imaging using statistical shape models. *Medical Image Analysis* **15** (2011)
5. Balakrishnan, G., Zhao, A., Sabuncu, M., Guttag, J., Dalca, A.V.: Voxelmorph: A learning framework for deformable medical image registration. *IEEE Transactions on Medical Imaging* **38** (2019)
6. Biguri, A., Dosanjh, M., Hancock, S., Soleimani, M.: Tigre: a matlab-gpu toolbox for cbct image reconstruction. *Biomedical Physics & Engineering Express* **2** (2016)
7. Cafaro, A., Spinat, Q., Leroy, A., Maury, P., Munoz, A., Beldjoudi, G., Robert, C., Deutsch, E., Grégoire, V., Lepetit, V., Paragios, N.: X2vision: 3d ct reconstruction from biplanar x-rays with deep structure prior. In: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.)

- International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI) (2023)
8. Cai, Y., Liang, Y., Wang, J., Wang, A., Zhang, Y., Yang, X., Zhou, Z., Yuille, A.: Radiative gaussian splatting for efficient x-ray novel view synthesis. In: European Conference on Computer Vision (ECCV) (2024)
 9. Cai, Y., Wang, J., Yuille, A., Zhou, Z., Wang, A.: Structure-aware sparse-view x-ray 3d reconstruction. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
 10. Cao, C.F., Ma, K.L., Shan, H., Liu, T.F., Zhao, S.Q., Wan, Y., Zhang, J., Wang, H.Q.: Ct scans and cancer risks: a systematic review and dose-response meta-analysis. *BMC Cancer* **22** (2022)
 11. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models-their training and application. *Computer Vision and Image Understanding* **61** (1995)
 12. D’Amore, T., Klein, G., Lonner, J.: The use of computerized tomography scans in elective knee and hip arthroplasty—what do they tell us and at what risk? *Arthroplasty Today* **15** (2022)
 13. Edgar, H.J.H., Daneshvari Berry, S., Moes, E., Adolphi, N.L., Bridges, P., Nolte, K.B.: New Mexico Decedent Image Database. Office of the Medical Investigator, University of New Mexico (2020)
 14. Feldkamp, L.A., Davis, L.C., Kress, J.W.: Practical cone-beam algorithm. *J. Opt. Soc. Am. A* **1** (1984)
 15. Guédon, A., Lepetit, V.: Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
 16. Heimann, T., Meinzer, H.P.: Statistical shape models for 3d medical image segmentation: a review. *Medical Image Analysis* **13** (2009)
 17. Hosseinian, S., Arefi, H.: 3d reconstruction from multi-view medical x-ray images—review and evaluation of existing methods. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **XL-1/W5** (2015)
 18. Kak, A.C., Slaney, M.: Principles of computerized tomographic imaging. SIAM (2001)
 19. Karade, V., Ravi, B.: 3d femur model reconstruction from biplane x-ray images: a novel method based on laplacian surface deformation. *International Journal of Computer Assisted Radiology and Surgery* **10** (2015)
 20. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42** (2023)
 21. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: International Conference on Learning Representations (ICLR) (2015)
 22. Qian, S., Kirschstein, T., Schoneveld, L., Davoli, D., Giebenhain, S., Nießner, M.: Gaussianavatars: Photorealistic head avatars with rigged 3d gaussians. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
 23. Rambani, R., Varghese, M.: Computer assisted navigation in orthopaedics and trauma surgery. *Orthopaedics and Trauma* **28** (2014)
 24. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena* **60** (1992)
 25. Turk, M., Pentland, A.: Eigenfaces for recognition. *Journal of Cognitive Neuroscience* **3** (1991)
 26. Tustison, N.J., Cook, P.A., Holbrook, A.J., Johnson, H.J., Muschelli, J., Devenyi, G.A., Duda, J.T., Das, S.R., Cullen, N.C., Gillen, D.L., Yassa, M.A., Stone, J.R.,

- Gee, J.C., Avants, B.B.: The ANTsX ecosystem for quantitative biological and medical imaging. *Scientific Reports* **11** (2021)
27. Van Houtte, J., Audenaert, E., Zheng, G., Sijbers, J.: Deep learning-based 2d/3d registration of an atlas to biplanar x-ray images. *International Journal of Computer Assisted Radiology and Surgery* **17** (2022)
 28. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **13** (2004)
 29. Wasserthal, J., Breit, H.C., Meyer, M.T., Pradella, M., Hinck, D., Sauter, A.W., Heye, T., Boll, D.T., Cyriac, J., Yang, S., Bach, M., Segeroth, M.: TotalSegmentator: Robust segmentation of 104 anatomic structures in CT images. *Radiology: Artificial Intelligence* **5** (2023)
 30. Zha, R., Lin, T.J., Cai, Y., Cao, J., Zhang, Y., Li, H.: R²-gaussian: Rectifying radiative gaussian splatting for tomographic reconstruction. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2024)
 31. Zha, R., Zhang, Y., Li, H.: Naf: neural attenuation fields for sparse-view cbct reconstruction. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (2022)
 32. Zhao, Z., Bao, Z., Li, Q., Qiu, G., Liu, K.: Psavatar: A point-based shape model for real-time head avatar animation with 3d gaussian splatting. *IEEE Transactions on Visualization and Computer Graphics* **32** (2026)