

ClipGS: Clippable Gaussian Splatting for Interactive Cinematic Visualization of Volumetric Medical Data

Chengkun Li¹, Yuqi Tong¹, Kai Chen¹, Zhenya Yang¹, Ruiyang Li¹, Shi Qiu¹, Jason Ying-Kuen Chan², Pheng-Ann Heng¹, and Qi Dou¹ (✉)

¹ Department of CSE, The Chinese University of Hong Kong

² Department of ENT, The Chinese University of Hong Kong

Abstract. The visualization of volumetric medical data is crucial for enhancing diagnostic accuracy and improving surgical planning and education. Cinematic rendering techniques significantly enrich this process by providing high-quality visualizations that convey intricate anatomical details, thereby facilitating better understanding and decision-making in medical contexts. However, the high computing cost and low rendering speed limit the requirement of interactive visualization in practical applications. In this paper, we introduce ClipGS, an innovative Gaussian splatting framework with the clipping plane supported, for interactive cinematic visualization of volumetric medical data. To address the challenges posed by dynamic interactions, we propose a learnable truncation scheme that automatically adjusts the visibility of Gaussian primitives in response to the clipping plane. Besides, we also design an adaptive adjustment model to dynamically adjust the deformation of Gaussians and refine the rendering performance. We validate our method on five volumetric medical data (including CT and anatomical slice data), and reach an average 36.635 PSNR rendering quality with 156 FPS and 16.1MB model size, outperforming state-of-the-art methods in rendering quality and efficiency. Project is available at: <https://med-air.github.io/ClipGS>.

Keywords: Volumetric Medical Data Visualization · Interactive Cinematic Rendering · Clippable Gaussian Splatting.

1 Introduction

Visualization of volumetric medical data [30] is a cornerstone of modern medical practice. The effectiveness of medical volume visualization is influenced by two key factors. First, the quality of rendering directly impacts the richness of the information conveyed. Recently, high-quality cinematic rendering technology [10] is widely used in surgical planning [25], medical education [12]. Second, the level of interactivity affects the user’s understanding of the data. It enables the intuitive visualization of complex medical data while effectively highlighting anatomical abnormalities, facilitated by interactive tools such as clipping planes [3,8].

To realize cinematic-level rendering of complex anatomical structures, ray tracing [15] is widely used to simulate the light propagation between voxels, enhancing detail and achieving high-quality photorealism [5,27]. To extend the usability of cinematic volume medical data visualization in practical applications, researchers [3,8,9] introduce clipping plane and clipping cube to interactively show the internal anatomy structures. However, the time-consuming rendering limits its practical application. To accelerate rendering processing, photon mapping-based volume rendering [29] is proposed in volumetric medical data visualization. Despite its advantages, the computational demands of this method still restrict real-time performance on consumer-grade hardware. Although Heinrich [11] and Taibo [26] further optimize the efficiency of interactive 3D medical visualization, their methods have to compromise rendering quality to achieve real-time interactions with volume clipping. It remains an open challenge to design a more efficient solution, implementing photo-realistic visualization while simultaneously supporting real-time interaction.

Recent advancements in neural rendering techniques have achieved great success in photo-realistic rendering with novel views. Neural radiance fields (NeRF) [19] represent 3D scenes as a continuous volumetric function through a multi-layer perception (MLP). More recently, 3D Gaussian Splatting (3DGS) models the 3D scene as a set of anisotropic Gaussian primitives [13,23,24,32] to implement the point-based radiance rendering process, achieving high-quality rendering and achieves impressive real-time performance. Benefiting from these, Niedermayr et al. [20] explored the potential capacity of 3DGS [13] technology for anatomical visualization. Kleinbeck et al. [14] took this a step further by extending a layered GS representation to visualize different anatomical structures incrementally. Although they implement real-time cinematic rendering of volumetric medical data, these methods only support rendering the fixed 3D structure and lack the ability to show internal structures and details. The non-interactive rendering limits the potential in surgical planning and education.

In this paper, we aim to establish a framework for real-time interactive cinematic visualization of volumetric medical data based on Gaussian Splatting. Here, we leverage the clipping plane, which is widely used in medical practice [3,8,9], to interactively visualize the internal structures and highlight the details. However, the use of a clipping plane introduces an additional variable, making it difficult to directly adapt to 3DGS-based rendering. Unlike existing method [14] that constructs 3D Gaussians for each clipping layer and exponentially increases storage/memory cost, we propose a novel GS-based framework that can render volumetric medical data with photorealism and clipping plane interaction in real-time. We introduce a learnable attribute to the Gaussian primitive to automatically control the visibility with the query clipping plane. Besides, we suggest dynamically adjusting the position and shape of visible Gaussian under the specific clipping plane to refine the rendering performance in the clipping surface. Our main contributions are summarized as follows:

- We propose ClipGS, a novel GS-based framework for real-time interactive cinematic anatomy visualization.




Fig. 1: Overview of our proposed framework. Our ClipGS implements a real-time interactive cinematic rendering under the query condition (view direction & clipping plane) for a given volumetric medical data.

- We propose a learnable truncation (LT) scheme, which can dynamically control the visibility of each Gaussian primitive with the query clipping plane.
- We design an adaptive adjustment model (AAM) to adjust the deformation of Gaussian to refine the rendering performance.
- We build a cinematic medical dataset to validate our proposed method, demonstrating its superiority in visual quality and efficiency.

2 Method

The overview pipeline of our proposed framework is shown in Fig. 1. For a given volumetric medical data $\mathbb{V}$, our method implements a real-time cinematic rendering for a given query view and clipping plane. Firstly, we use a ray-tracing renderer to generate a sparse cinematic image sequence $\{\mathbf{I}_i\}_{i=1}^N$ under random camera views $\mathbf{v}_i$ and random clipping plane $\mathbf{z}_i$. Then, we optimize a random initialized 3D Gaussian point cloud (Sec. 2.1) from these images. For doing this, We introduce a learnable attribute to the Gaussian primitive to automatically control the visibility with the query clipping plane (Sec. 2.2). To refine the rendering quality, we design an adaptive adjustment model to dynamically adjust the position and shape of visible Gaussian under the specific clipping plane (Sec. 2.3). Our experiments show the superiority of the proposed method in rendering quality and efficiency (Sec. 3).

2.1 Preliminaries of Gaussian Splatting

3D Gaussian splatting [13] represents the static 3D scene with a set of Gaussian primitives. Mathematically, it could be written as:

$$\mathcal{G}(\mathbf{x}; \boldsymbol{\mu}, \mathbf{r}, s) = e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})}, \quad \boldsymbol{\Sigma} = \mathbf{R}\mathbf{S}\mathbf{S}^T\mathbf{R}^T, \quad (1)$$

where the covariance matrix is decomposed as the rotation $\mathbf{R}$ and scaling $\mathbf{S}$. For the rendering process, 3DGS uses a differentiable rasterizer to implement point-based volume rendering. The pixel color $\mathbf{C}$ in the image plane is finally rendered by the α -blending of the projected Gaussian primitive on this plane.

$$\mathbf{C} = \sum_{i=1}^N T_i \alpha_i \mathcal{G}'_i(\mathbf{x}; \boldsymbol{\mu}_i, \mathbf{r}_i, s_i) \mathbf{c}_i, \quad T_i = \prod_{j \in i-1} (1 - \mathcal{G}'_j(\mathbf{x}; \boldsymbol{\mu}_j, \mathbf{r}_j, s_j) \alpha_j), \quad (2)$$

where T_i indicates the accumulated transmittance of i -th Gaussian at pixel $\mathbf{x}$. The projected 2D Gaussian probability density $\mathcal{G}'(\mathbf{x}; \boldsymbol{\mu}, \mathbf{r}, s)$ could be obtained by projection operation from the camera's intrinsic and extrinsic parameters. Finally, the optimizable Gaussian primitives are characterized by: position $\boldsymbol{\mu} \in \mathbb{R}^3$, rotation $\mathbf{r} \in \mathbb{R}^3$, scale factor $s \in \mathbb{R}$, color $\mathbf{c} \in \mathbb{R}^3$ and opacity $\alpha \in \mathbb{R}$.

2.2 Gaussian Visibility Control via Learnable Truncation

The clipping plane factor introduces an additional dimension for cinematic rendering baking of 3D medical data with Gaussian splatting, which transfers the 3D representation to the 4D representation problem. Previous works [7, 18, 28] try to learn a deformation field along the introduced dimension to control the rendering of GS. But, these methods conflict with the inspiration that the applied clipping plane mainly changes the visibility of Gaussian primitives during rendering. Inspired by the above, we consider the geometric relationship between the clipping plane and Gaussian primitives.

Intuitively, we can directly use the mean value $\boldsymbol{\mu}$ as the position of Gaussian to judge the visibility, which we refer to as the Hard Truncation (HT) scheme. Although the rendering results are guaranteed to be generally correct, this simple strategy tends to trap models in local minima, causing artifacts to burr in the clipping surface. In a specific case, we consider the Gaussian primitives, which have an intersection with the clipping plane under 99% confidence interval as per 3σ principle. They contribute color to the rendering of the clipping surface, but their mean value may fall on both sides of the clipping plane. Instead, we consider the real ‘‘barycenter’’ of Gaussian actually has an offset $\boldsymbol{\delta} \in \mathbb{R}^3$ respects to mean value $\boldsymbol{\mu}$.

$$\mathcal{M} = (\boldsymbol{\mu} + \boldsymbol{\delta}) \cdot \mathbf{n} < z, \quad (3)$$

where $\mathbf{n}$ is the normal direction of the clipping plane, and z is the distance of the clipping plane to the origin point. $\mathcal{M} \in \{0, 1\}$ means the visibility. Gaussians under the clipping plane are visible, and vice versa.

To this end, we propose a Learnable Truncation (LT) scheme by introducing an additional optimizable attribute $m \in \mathbb{R}$ to each Gaussian primitive. We formulate $m = (\boldsymbol{\mu} + \boldsymbol{\delta}) \cdot \mathbf{n}$, and rewrite Eq. 3 as $\mathcal{M} = m < z$. Thus, we are able to ensure that the operation of $\boldsymbol{\mu}$ is not affected by the gradient of Eq. 3. Inspired by [16], we employ the straight-through estimator [2], and redefine our learnable truncation function as:

$$\mathcal{M} = \text{sg}(\mathbb{1}[\sigma(m - z) < \epsilon] - \sigma(m - z)) + \sigma(m - z), \quad (4)$$

where $\sigma(\cdot)$ is the sigmoid function, $\mathbb{1}(\cdot)$ is an indicator, ϵ is a threshold hyperparameter and $\text{sg}(\cdot)$ is the stop gradient operator. In this way, we can directly use the gradient of the sigmoid function to represent the gradient of the visibility step function during backpropagation, and optimizing the value of m automatically.

2.3 Continuous Clipping via Adaptive Deformation Adjustment

Compared to the 3D scene representation task, interactive cinematic visualization introduces an additional variable to show the internal structures of the volumetric medical data. The introduced dimension transfers this task as a 4D scene representation problem. But existing GS-based cinematic rendering [20,14] only addresses 3D problems. We use a novel learnable truncation scheme (Sec. 2.2) to control the visibility of Gaussian primitives for cinematic anatomy rendering. However, it is mathematically non-contiguous due to the indicator operation in Eq. 4, causing inconsistent rendering when the clipping plane changes. Instead of performing a direct visibility control, we propose to predict a continuous adjustment to the Gaussian primitives based on the clipping plane condition, effectively refining the rendering results and ensuring continuity across clipping planes.

To deal with this issue, we design an adaptive adjustment model to model the dynamic deformation of Gaussians under the specific clipping plane condition. We exploit positional encoding (PE) [19] followed by a tiny MLP to learn a feature f , and the desired dynamic deformations are generated from a multi-head MLP decoder. Here, we respectively input all visible Gaussian positions $\boldsymbol{\mu}$ and clipping plane parameter z into the PE, and then the merged high-frequency input from PE is fed into the MLP.

$$f = \text{MLP}(\gamma(\boldsymbol{\mu}), \gamma(z)), \quad (5)$$

where $\gamma(\cdot)$ denotes the positional encoding operation. Then separate MLP heads ϕ_μ , ϕ_γ , and ϕ_s , are employed to compute the deformations of Gaussians: position $\Delta\boldsymbol{\mu} = \phi_\mu(f)$, rotation $\Delta\mathbf{r} = \phi_r(f)$ and scale $\Delta\boldsymbol{\mu} = \phi_s(f)$. Finally, we rewrite the rendering function in Eq. 2 as follows:

$$\mathcal{C} = \sum_{i=1}^N \mathcal{M}_i T_i \alpha_i \mathcal{G}'_i(\mathbf{x}; \boldsymbol{\mu}_i + \Delta\boldsymbol{\mu}_i, \mathbf{r}_i + \Delta\mathbf{r}_i, s_i + \Delta s_i) \mathbf{c}_i. \quad (6)$$

The rendered result is colored by α -blending of all visible Gaussians with the query clipping plane.

2.4 Optimization of ClipGS

We employ a two-step optimization strategy in our method. Firstly, we train the GS with our learnable truncation scheme to obtain a coarse initialization. Then, we start the adaptive adjustment to refine the rendering performance and keep the consistency in the clipping plane dimension. As for the training loss, we measure the difference between the rendered colors $\mathcal{C}$ and the ground-truth

colors $\mathbf{C}_{\text{gt}}$ from the rendered CA image. Similar as 3DGS pipeline in [13], we use the combination of the D-SSIM term and a $\mathcal{L}_1$ term as our training loss:

$$\mathcal{L} = (1 - \lambda)\mathcal{L}_1(\mathbf{C}, \mathbf{C}_{\text{gt}}) + \lambda\mathcal{L}_{\text{D-SSIM}}(\mathbf{C}, \mathbf{C}_{\text{gt}}). \quad (7)$$

In this way, we are able to balance the perceptual quality of the rendered images with the fidelity to the real data.

3 Experiments

3.1 Experiment Settings

Dataset. To evaluate the rendering performance of our proposed method on volume medical data, we generate a custom cinematic medical dataset by BioxelNode [21] of Blender [6]. The custom cinematic medical dataset consists of five scenes: 1) *Head*; 2) *Lower Limb*; 3) *Skull*; 4) *Thorax*; 5) *Liver & Kidney*. The original volume of medical data (scenes 1st to 4th) was acquired from the Visible Human Project [1], and the data of scene 5th was acquired from CVH project [31]. We employ a DyNeRF-like [17] format in our dataset, which contains a total of 6000 frames with different camera views and timestamps. We rendered 1100 frames for each scene, with random view direction on a unit sphere and random clipping plane position along the z-axis for each frame, all at 800×800 pixels. We use 1000 frames for training and 100 frames for testing.

Implementation Details. For each scene, we randomly initialize the GS with 100,000 points within a cube encompassing the scene as default. We train 7000 iterations in the first optimization step, and jointly optimize the GS with our AAM for another 33000 iterations. We employ $\epsilon = 0.5$ in Eq. 4 and $\lambda = 0.2$ in Eq. 7. For other hyperparameters, we use the same with 3DGS. We implement our method based on the PyTorch framework [22] and conduct the experiments with a single RTX 4090 GPU.

3.2 Qualitative and Quantitative Results

How does our method compare to others? The clipping plane view introduces an additional dimension for the NVS task. Actually, it could be seen as a 4D scene representation problem. Similarly, the clipping plane can be considered as the timestamp in the dynamic scene representation task. Thus, we select HexPlane [4], GauFRe [18], N-DG [7] and 4DGaussians [28] for comparative. To assess our method’s performance, we used Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) and perceptual quality measure (LPIPS) for image quality, model size for spatial efficiency, and Frames Per Second (FPS) for rendering speed.

What is our model’s performance, and why it works? Our proposed model gives optimal rendering performance among all methods, and also has




Fig. 2: Qualitative results of rendered images on our cinematic medical dataset. The PSNRs in the images denote the rendering quality for the whole case.

great performance in training time, reducing storage and also rendering speed, as shown in Tab. 1 and Fig. 2. Other compared methods aim to learn a deformation field to represent the rendering variance along the additional introduced dimension (clipping plane). A significant drawback arises as these methods entail a large model capacity to explain this variance. Instead, we use a novel learnable truncation to incorporate the space geometric relationship between the Gaussian primitives and the clipping plane. Our model allows for a main focus on the visible Gaussian primitive optimization under a given clipping plane condition, additional AAM to refine the rendering result.

3.3 Ablation Study

Learnable Truncation Scheme. The proposed learnable truncation allows our model to focus mainly on the visible Gaussian primitive optimization under a given clipping plane condition. Compared to the quantitative results (3DGS + AAM vs. 3DGS + AAM + LT) in Tab. 2, our LT scheme shows efficient performance in reducing the floaters, training time, saving storage, and improving

Table 1: Quantitative performance. We compared the results in photometric loss, and also the performance in training time (min), rendering speed, and saving storage on the cinematic anatomy dataset.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Time $\downarrow$	FPS $\uparrow$	Storage $\downarrow$
HexPlane [4]	31.836 ± 1.275	0.942 ± 0.015	0.109 ± 0.019	14.1'	2.7	68.0MB
GauFRe [18]	34.717 ± 1.481	0.968 ± 0.009	0.062 ± 0.013	16.6'	150	15.0MB
N-DG [7]	31.925 ± 1.541	0.937 ± 0.032	0.096 ± 0.015	33.6'	129	39.7MB
4DGaussians [28]	35.653 ± 1.667	0.967 ± 0.009	0.079 ± 0.014	9.7'	87	17.8MB
Ours	36.635 ± 1.926	0.974 ± 0.010	0.061 ± 0.014	13.7'	156	16.1MB

Table 2: Ablation study of the proposed components. We compared the effectiveness of our learnable truncation (LT) and adaptive adjustment model (AAM). Also, we show the superiority of our method compared to the hard truncation (HT) scheme. 3DGS + AAM + LT represents our full model.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Time $\downarrow$	FPS $\uparrow$	Storage $\downarrow$
3DGS + LT	34.577 ± 1.441	0.969 ± 0.012	0.070 ± 0.015	11.1'	201	14.7MB
3DGS + AAM	35.946 ± 2.043	0.971 ± 0.013	0.064 ± 0.014	11.5'	148	15.6MB
3DGS + AAM + HT	36.255 ± 1.756	0.971 ± 0.008	0.063 ± 0.014	14.1'	150	15.7MB
3DGS + AAM + LT	36.635 ± 1.926	0.974 ± 0.010	0.061 ± 0.014	13.7'	156	16.1MB

details. Besides, we also compared our Learnable truncation with the Hard truncation, which has been mentioned in Sec. 2.2 (3DGS + AAM + LT vs. 3DGS + AAM + HT); HT has almost the same performance in reducing training time and saving the storage. However, the non-optimizable strategy compromises the rendering quality of the model.

Adaptive Adjustment Model. Our adaptive adjustment model predicts a dynamic deformation for the Gaussian. So that the position, rotation and scaling of these Gaussian primitives could be adaptively adjusted given the specific clipping plane. By incorporating an additional MLP to learn the clipping plane-dependent variations, our proposed module exhibits greater capacity to mitigate photometric loss, thereby achieving improved results, as shown in Table 2 (3DGS + LT vs. 3DGS + AAM + LT).

4 Conclusion

In this paper, we proposed a novel GS-based framework that can render volumetric medical data with clipping planes with photorealism and real-time. We propose a learnable truncation scheme that automatically adjusts the visibility of Gaussian primitives in response to the clipping plane. In addition, we propose an adaptive adjustment model to model the dynamic deformation of Gaussians under the specific clipping plane condition. The experiments on our cinematic

medical dataset demonstrate that our method achieves superior image quality, rendering speed, and storage efficiency compared to existing approaches, making it a promising solution for more accessible and interactive medical imaging applications in surgical planning and medical education.

Acknowledgments. This research work was supported in part by the Research Grants Council of the Hong Kong Special Administrative Region, China, under Projects No. 24209223 and No. T45-401/22-N, in part by the National Natural Science Foundation of China under Project No. 62322318, in part by the Hong Kong Innovation and Technology Fund under Project No. GHP/167/22SZ, and in part by the Chinese University of Hong Kong under Project No. 4055212.

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

References

1. Ackerman, M.: The visible human project. *Proceedings of the IEEE* **86**(3), 504–511 (1998). <https://doi.org/10.1109/5.662875> 6
2. Bengio, Y., Léonard, N., Courville, A.: Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432* (2013) 4
3. Binder, J., Krautz, C., Engel, K., Grützmann, R., Fellner, F.A., Burger, P.H., Scholz, M.: Leveraging medical imaging for medical education—a cinematic rendering-featured lecture. *Annals of Anatomy-Anatomischer Anzeiger* **222**, 159–165 (2019) 1, 2
4. Cao, A., Johnson, J.: Hexplane: A fast representation for dynamic scenes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 130–141 (2023) 6, 8
5. Comaniciu, D., Engel, K., Georgescu, B., Mansi, T.: Shaping the future through innovations: From medical imaging to precision medicine (2016) 2
6. Community, B.O.: Blender - a 3D modelling and rendering package. Blender Foundation, Stichting Blender Foundation, Amsterdam (2018), <http://www.blender.org> 6
7. Diolatzis, S., Zirr, T., Kuznetsov, A., Kopanas, G., Kaplanyan, A.: N-dimensional gaussians for fitting of high dimensional functions. In: *ACM SIGGRAPH 2024 Conference Papers*. pp. 1–11 (2024) 4, 6, 8
8. Eid, M., De Cecco, C.N., Nance Jr, J.W., Caruso, D., Albrecht, M.H., Spandorfer, A.J., De Santis, D., Varga-Szemes, A., Schoepf, U.J.: Cinematic rendering in ct: a novel, lifelike 3d visualization technique. *American Journal of Roentgenology* **209**(2), 370–379 (2017) 1, 2
9. Elshafei, M., Binder, J., Baecker, J., Brunner, M., Uder, M., Weber, G.F., Grützmann, R., Krautz, C.: Comparison of cinematic rendering and computed tomography for speed and comprehension of surgical anatomy. *JAMA surgery* **154**(8), 738–744 (2019) 2
10. Glemser, P.A., Engel, K., Simons, D., Steffens, J., Schlemmer, H.P., Orakcioglu, B.: A new approach for photorealistic visualization of rendered computed tomography images. *World neurosurgery* **114**, e283–e292 (2018) 1

11. Heinrich, F., Bornemann, K., Lawonn, K., Hansen, C.: Interacting with medical volume data in projective augmented reality. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 429–439. Springer (2020) [2](#)
12. Jabbireddy, S., Li, S., Meng, X., Terrill, J.E., Varshney, A.: Accelerated volume rendering with volume guided neural denoising. In: EuroVis (Short Papers). pp. 49–53 (2023) [1](#)
13. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42**(4) (July 2023), <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/> [2](#), [3](#), [6](#)
14. Kleinbeck, C., Schieber, H., Engel, K., Gutjahr, R., Roth, D.: Multi-layer gaussian splatting for immersive anatomy visualization. *arXiv preprint arXiv:2410.16978* (2024) [2](#), [5](#)
15. Lafortune, E.P., Willems, Y.D.: Bi-directional path tracing (1993) [2](#)
16. Lee, J.C., Rho, D., Sun, X., Ko, J.H., Park, E.: Compact 3d gaussian representation for radiance field. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21719–21728 (2024) [4](#)
17. Li, T., Slavcheva, M., Zollhoefer, M., Green, S., Lassner, C., Kim, C., Schmidt, T., Lovegrove, S., Goesele, M., Newcombe, R., et al.: Neural 3d video synthesis from multi-view video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5521–5531 (2022) [6](#)
18. Liang, Y., Khan, N., Li, Z., Nguyen-Phuoc, T., Lanman, D., Tompkin, J., Xiao, L.: Gafre: Gaussian deformation fields for real-time dynamic novel view synthesis. In: WACV (2025) [4](#), [6](#), [8](#)
19. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021) [2](#), [5](#)
20. Niedermayr, S., Neuhauser, C., Petkov, K., Engel, K., Westermann, R.: Application of 3d gaussian splatting for cinematic anatomy on consumer class devices. *Vision, Modeling, and Visualization*, L. Linsen and J. Thies, Eds. The Eurographics Association (2024) [2](#), [5](#)
21. OmooLab: Bioxelnodes (2025), <https://github.com/OmooLab/BioxelNodes>, accessed: 2025-02-27 [6](#)
22. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32** (2019) [6](#)
23. Qiu, S., Xie, B., Liu, Q., Heng, P.A.: Advancing extended reality with 3d gaussian splatting: Innovations and prospects. In: 2025 IEEE International Conference on Artificial Intelligence and eXtended and Virtual Reality (AIxVR). pp. 203–208 (2025) [2](#)
24. Qiu, S., Xie, B., Liu, Q., Heng, P.A.: Creating virtual environments with 3d gaussian splatting: A comparative study. In: 2025 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW). pp. 1332–1333 (2025) [2](#)
25. Steffen, T., Winklhofer, S., Starz, F., Wiedemeier, D., Ahmadli, U., Stadlinger, B.: Three-dimensional perception of cinematic rendering versus conventional volume rendering using ct and cbct data of the facial skeleton. *Annals of Anatomy-Anatomischer Anzeiger* **241**, 151905 (2022) [1](#)
26. Taibo, J., Iglesias-Guitian, J.A.: Immersive 3d medical visualization in virtual reality using stereoscopic volumetric path tracing. In: 2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR). pp. 1044–1053. IEEE (2024) [2](#)

27. Veach, E., Guibas, L.: Bidirectional estimators for light transport. In: Photorealistic Rendering Techniques, pp. 145–167. Springer (1995) [2](#)
28. Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Wang, X.: 4d gaussian splatting for real-time dynamic scene rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20310–20320 (2024) [4](#), [6](#), [8](#)
29. Yuan, Y., Yang, J., Sun, Q., Huang, Y., Ma, S.: Cinematic volume rendering algorithm based on multiple lights photon mapping. Multimedia Tools and Applications **83**(2), 5799–5812 (2024) [2](#)
30. Zhang, Q., Eagleson, R., Peters, T.M.: Volume visualization: a technical overview with a focus on medical applications. Journal of digital imaging **24**, 640–664 (2011) [1](#)
31. Zhang, S.X., Heng, P.A., Liu, Z.J.: Chinese visible human project: Dataset acquisition and its primary applications. In: 2005 IEEE Engineering in Medicine and Biology 27th Annual Conference. pp. 4168–4170 (2005) [6](#)
32. Zhu, R., Qiu, S., Liu, Z., Hui, K.H., Wu, Q., Heng, P.A., Fu, C.W.: Rethinking end-to-end 2d to 3d scene segmentation in gaussian splatting. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 3656–3665 (2025) [2](#)