

Article

Not peer-reviewed version

Efficient Multi-Material Volume Rendering for Realistic Visualization with Complex Transfer Functions

[Chunxiao Xu](#), [Xinran Xu](#), [Jiatian Zhang](#), [Yiheng Cao](#), [Lingxiao Zhao](#)*

Posted Date: 13 May 2025

doi: 10.20944/preprints202505.0946.v1

Keywords: Medical visualization; Realistic volume rendering; Volumetric Accelerator; Transfer function








Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Efficient Multi-Material Volume Rendering for Realistic Visualization with Complex Transfer Functions

Chunxiao Xu ^{1,2} , Xinran Xu ^{1,2} , Jiatian Zhang ^{1,2} , Yiheng Cao ²  and Lingxiao Zhao ^{1,2,*} 

¹ School of Biomedical Engineering, Division of Life Sciences and Medicine, University of Science and Technology of China, Hefei, Anhui, 230026, P.R.China

² Suzhou Institute of Biomedical Engineering and Technology, Chinese Academy of Science, Suzhou, Jiangsu, 215613, P.R.China

* Correspondence: hitic@sibet.ac.cn

Abstract: Physically-based realistic direct volume rendering (DVR) is a critical area of research in scientific data visualization. The prevailing realistic DVR methods are primarily rooted in outdated theories of participating media rendering and often lack comprehensive analyses of their applicability to realistic DVR scenarios. As a result, the fidelity of material representation in the rendered output is frequently limited. To address these challenges, we present a novel multi-material radiative transfer model (MM-RTM) designed for realistic DVR, grounded in recent advancements in light transport theories. Additionally, we standardize various transfer function techniques and propose five distinct forms of transfer functions along with proxy volumes. This comprehensive approach enables our DVR framework to accommodate a wide range of complex transfer function techniques, which we illustrate through several visualizations. Furthermore, to enhance sampling efficiency, we develop a new multi-hierarchical volumetric acceleration method that supports multi-level searches and volume traversal. Our volumetric accelerator also facilitates real-time structural updates when applying complex transfer functions in DVR. Our MM-RTM, the unified representation of complex transfer functions, and the acceleration structure for real-time updates are complementary components that collectively establish a comprehensive framework for realistic multi-material DVR. Evaluation from a user study indicates that the rendering results produced by our method demonstrate the most realistic effects among various publicly available state-of-the-art techniques.

Keywords: medical visualization; realistic volume rendering; volumetric accelerator; transfer function

1. Introduction

Direct Volume Rendering (DVR) plays an indispensable role in scientific data visualization [1–3], especially for the informative presentation and interpretation of medical image data [1,4]. DVR techniques [5,6] primarily evolved from simulations of light transport [7–9]. They make use of transfer functions [10] to map volumetric data values to material properties. With the development of computing technology, specifically benefitting from the significantly improved performance of Graphic Processing Units (GPUs) and the refined DVR theories, researchers have been able to explore the application of tailored DVR techniques for specified visualization tasks and effects [11–13].

The realism of rendering comes from simulating light transport and how light interacts with materials [14]. Some research works have verified that incorporating advanced illumination techniques in DVR can create realistic representations of materials with shadow casting and help enhance the observer's spatial awareness and comprehension of the data [15,16]. Numerous studies have attempted to integrate various illumination algorithms in DVR to produce different shading effects and visual outcomes [16–19]. The volumetric path tracing (VPT) based on Monte Carlo (MC) estimations has been introduced in DVR to generate photorealistic rendering results [20–22]. VPT-based cinematic

rendering techniques have been predominantly utilized in the 3D visualization of medical images, demonstrating substantial usability in the clinical diagnosis of diseases [4,23–25].

The radiative transfer equation (RTE) can be used to simulate the light propagation within non-refractive participating media. VPT-based methods for estimating RTE developed in recent years [26–29] were primarily based on the null-collision technique [30]. These methods deserve attention for their conciseness, efficiency and analytical tractability. Although many studies advocate the use of VPT as the DVR technique, they often lack specific implementation details. Several implementations are derived from the ray-marching method used in ExposureRenderer [31]. In practice, using VPT for DVR without modifications usually produces visualizations with low contrast and insufficient detail. We analyze this issue and propose effective solutions to enhance the realism of VPT-based DVR.

The transfer function maps voxel values to various optical properties [10]. The diverse transfer function techniques employed by existing DVR methods complicate their integration into one system. To address this, we analyzed and unified the various transfer function techniques, proposing five primary types. We demonstrate that current transfer function techniques can be transformed into the following form: N (where $N \geq 1$) proxy volumes with N -dimensional transfer functions of these five types. Our classification approach enhances the utilization of multi-material radiative transfer models (MM-RTM), improving the realism of volume rendering results.

When the density within the volumetric space exhibits significant variation, a volume acceleration structure is typically employed to enhance sampling efficiency [32]. However, existing acceleration methods often encounter issues related to sampling efficiency or interactivity. Interactive DVR requires the acceleration structure to respond in real time to changes in the transfer function while ensuring that the dynamically constructed structure aligns closely with the volumetric density. To address this, a hierarchical volume acceleration structure is designed that supports empty space skips and multi-level volume traversal [14], significantly improving rendering efficiency compared to existing methods.

Our research work tries to extend the RTE to support realistic DVR for physically-based visualization of multiple materials. Realistic DVR results generated using our method are depicted in Figure 1, Unless otherwise specified, the datasets utilized in this paper are sourced from [33,34]. Our contributions include:

- We develop a new MM-RTM that supports the VPT-based DVR of multiple materials. Our method is capable of incorporating theoretical multiple scattering effects into realistic DVR.
- We propose a novel hierarchical volume grid acceleration structure. An update mechanism is designed for this acceleration structure that ensures accurate real-time updates during user interaction. Our update method supports applications involving complex transfer functions.
- We standardized the representation of transfer functions and identified five fundamental types. We demonstrate that complex transfer functions can be expressed in these five forms.

Our MM-RTM, the unified representation of complex transfer functions, and the acceleration structure for real-time updates are complementary components that together create a comprehensive framework for realistic DVR. We validated the effectiveness of the system through two medical applications, and after evaluation by medical experts, the system is adequate for surgical planning in these two surgical scenarios. In our supplementary material¹, we provide rendering results for real-time browsing. We include scenes reconstructed using neural rendering techniques, featuring real-time browsing effects achieved with neural radiance field (NeRF) [35,36] and 3D Gaussian splatting (3DGS) [37,38].

¹ [NeRF and 3DGS Data](#)

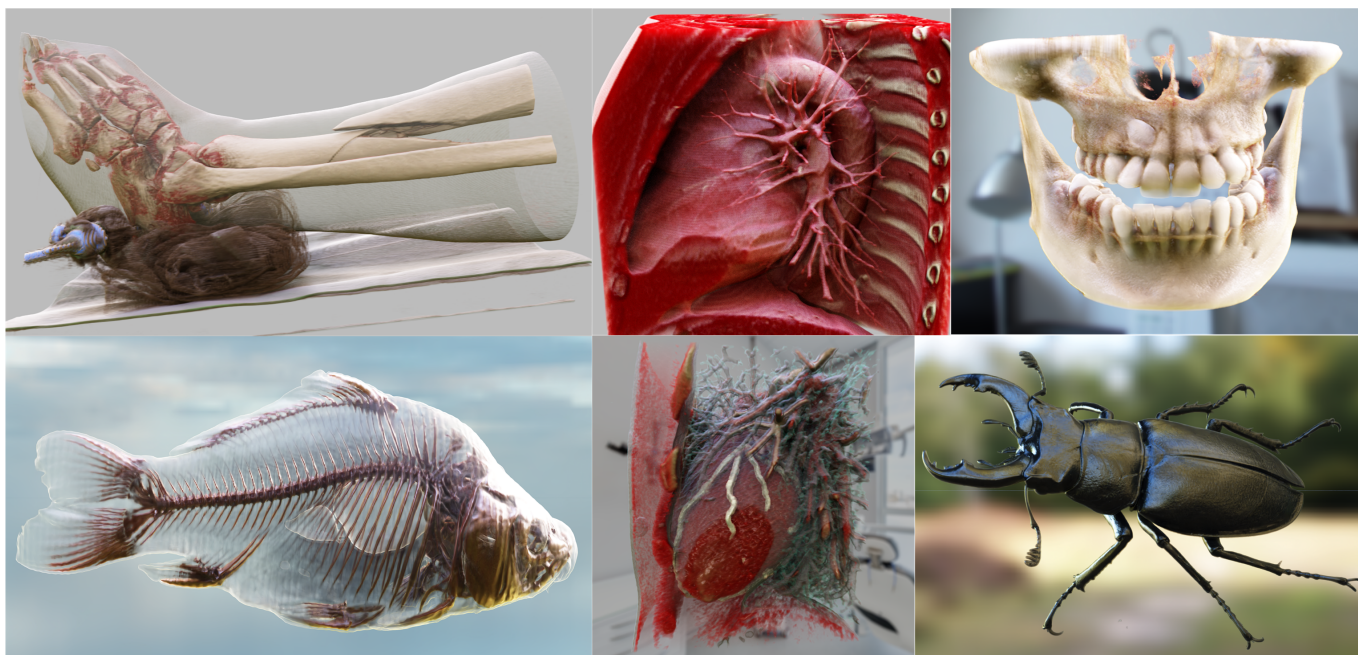


Figure 1. Our multi-material radiative transfer model supports physically-based rendering of multiple scattering effects. It can be used to significantly enhance the visual authenticity of the volume rendering results. Our rendering results exhibit high realism, such as the subsurface scattering effect observed in teeth, the smooth surface of the beetle and the semi-translucent appearance of muscles and blood vessels.

2. Related Work

The DVR technique has a wide range of applications in the field of scientific data visualization. It can serve to facilitate the human perception and understanding of volumetric data. Various literatures offer in-depth insights into this research domain [1,3]. Given that scientific data is usually derived from nature, visualization outcomes tend to possess a certain degree of realism or offer visual enhancements for specific regions of interest [10,12,15,39]. The realism of rendering often stems from the simulation of light transport and interactions between light and materials. [14,16].

Some research works have demonstrated that advanced illumination can enhance people's spatial perception of rendering results [1,15]. A comprehensive overview of advanced volumetric illumination techniques in interactive DVR is provided in the survey [16]. Numerous classic rendering techniques, including Precomputed Radiance Transfer (PRT), Photon Mapping (PM) and Virtual Point Lights (VPLs), have been adopted in DVR to achieve advanced shading effects. Bauer et al. [40] introduced photon field networks, which is a neural representation that facilitates interactive rendering with diminished stochastic noise. Wu et al. [41] proposed an efficient technique for volumetric neural representation. Their method performs real-time volumetric photon tracing using a newly developed sampling stream algorithm.

Advanced transfer function techniques can more effectively represent volumetric data information [10]. Various transfer function techniques use low-dimensional, high-dimensional and even semantic information for classification. Igouchkine et al. [12] introduced a material transfer function that facilitates high-quality rendering of multi-material volumes by generating surface-like behavior at structure boundaries and volume-like behavior within the structures. With the advancement of deep learning, automated classification methods based on semantic information have emerged. Nguyen et al. [42] proposed a transfer function technique utilizing deep learning to generate soft labels for visualizing electron microscopy data with high noise levels. Engel et al. [43] employed a large network model for semantic pre-classification of volumetric data, effectively distinguishing between different tissues and organs. Li et al. [44] proposed an attention-driven visual emphasis method for the visualization of volumetric medical images, focusing on the characterization of small regions of interest (ROIs).

In recent years, research efforts on rendering of participating media have made remarkable progress in both theories and applications [30,45]. VPT-based rendering of participating media involves techniques for sampling transmittance and free paths. Unbiased transmittance estimation includes ratio tracking and residual tracking [26]. Unbiased free path sampling techniques encompass decomposition tracking [27] and weighted tracking [30], as well as spectral tracing techniques [27] and null-scattering techniques [28,29] for estimating colored media through singlepath estimation. These techniques are mainly used for rendering sparse participating media such as fog, clouds, fire and milk-like fluids.

VPT and physically-based shading techniques have been widely used in DVR to generate visually realistic renderings [20–22]. Realistic DVR techniques primarily rely on incorporating surface material characteristics into volumetric shading [16]. Cinematic rendering approaches of medical images [4,25,46–48] have become a rapidly developing research direction. These techniques can be used to effectively assist doctors in understanding and analyzing pathologies. Denisova et al. [49] proposed a DVR system that accommodates multiple scattering. VPT techniques based on MC estimation often contain lots of noise during interaction. In recent years, denoising techniques for MC-based rendering have been introduced into VPT [50–53] to enhance the rendering quality.

3. Multi-Material Radiative Transfer Model

In this section we introduce our MM-RTM. The meanings of the symbols used in the equations are provided in Table 1. Assuming the concurrent presence of four types of particles with various optical properties in the volume, i.e., absorptive particles, volumetric scattering particles, virtual particles and surface scattering particles, the new characteristics of our MM-RTM can be modeled as follows:

- For particles exhibiting surface-scattering properties, we assume that the geometric normal is the volumetric gradient. The surface-scattering property characterizes hemispherical scattering. Incident light from the back side of the surface is fully absorbed by the surface-scattering particles.
- The absorptivity of a traditional medium can be elucidated as the absorption of light per unit distance [14]. For surface-scattering particles, absorptivity also encompasses the absorption of light energy per unit surface area.

It is indeed feasible to modify state-of-the-art bidirectional scattering distribution function (BSDF) to eliminate surface absorption and incorporate them into our MM-RTM. But the heuristic "scattering absorptivity" plays a crucial role in realistic DVR. It can strongly influence the appearance of rendering results. This obviates the necessity of employing multi-channel extinction coefficient σ_t for rendering chromatic volume.

Table 1. Notations of the RTE

Symbol	Description
\mathbf{p}	a point in volume space
$r(\mathbf{p}, \omega)$	a ray, emitted from point \mathbf{p} in the direction of ω
$L(\mathbf{p}, \omega)$	light from the point \mathbf{p} to the direction ω
$\sigma_a(\mathbf{p})$	absorption coefficient at the point \mathbf{p}
$\sigma_s(\mathbf{p})$	scattering coefficient at the point \mathbf{p}
$\sigma_t(\mathbf{p})$	extinction coefficient at the point \mathbf{p}
$\sigma_n(\mathbf{p})$	null-scattering coefficient at the point \mathbf{p}
σ_{maj}	the "majorant" constant, $\sigma_{\text{maj}} = \sigma_t(\mathbf{p}) + \sigma_n(\mathbf{p})$
$\int_{\mathcal{S}^2}$	spherical integral
$\int_{\mathcal{H}^2}$	the integral over the hemisphere
$\rho(\cdot)$	the phase function
$f(\cdot)$	the BSDF
$L_e(\mathbf{p}, \omega)$	radiance emitted to ω from \mathbf{p}
$T_r(\mathbf{p}_1 \leftrightarrow \mathbf{p}_2)$	light transmittance between \mathbf{p}_1 and \mathbf{p}_2 .
$\mathbf{G}(\mathbf{p})$	the gradient at the point \mathbf{p}
$A(\mathbf{p})$	the surface scattering albedo coefficient at \mathbf{p}
$S(\mathbf{p})$	the spherical scattering albedo coefficient at \mathbf{p}

3.1. Extension of RTE

The constraints for the phase function $\rho(\cdot)$ and the BSDF $f(\cdot)$ are different, primarily attributed to BSDF's inclusion of the albedo:

$$\int_{S^2} \rho(\omega_i, \mathbf{p}, \omega_o) d\omega_o = 1 \quad \int_{\mathcal{H}^2} f(\omega_i, \mathbf{p}, \omega_o) |\mathbf{G}(\mathbf{p}) \cdot \omega_o| d\omega_o \leq 1 \quad (1)$$

Denote albedo at \mathbf{p} as $A(\mathbf{p})$, the BSDF can be reformulated to derive the Proxy Scattering Function (PSF) f_s as:

$$f_s(\omega_i, \mathbf{p}, \omega_o) = \begin{cases} \frac{f(\omega_i, \mathbf{p}, \omega_o) |\mathbf{G}(\mathbf{p}) \cdot \omega_o|}{A(\mathbf{p})} & \mathbf{G}(\mathbf{p}) \cdot \omega_o > 0 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where \mathbf{G} is geometry item [14]. The integral of $f_s(\omega_i, \mathbf{p}, \omega_o)$ over the spherical surface is equal to 1.0.

In volume, there can be various surface scattering particles and volumetric scattering particles. For the sake of simplicity in formula derivation, it can be assumed that only one type of volumetric scattering particle (with phase function ρ) and one type of surface-scattering particle (with PSF f_s) exist simultaneously. The probabilities associated with each particle type are governed by the probability P , ensuring $P_{f_s}(\mathbf{p}) + P_\rho(\mathbf{p}) = 1$. Therefore, the RTE can be defined as:

$$L_n(\mathbf{p}, \omega) = \frac{\sigma_a(\mathbf{p})}{\sigma_{\text{maj}}} L_e(\mathbf{p}, \omega) + \frac{\sigma_n(\mathbf{p})}{\sigma_{\text{maj}}} L(\mathbf{p}, \omega) + \frac{P_\rho(\mathbf{p}) \sigma_s(\mathbf{p})}{\sigma_{\text{maj}}} \int_{S^2} \rho(\omega_i, \mathbf{p}, \omega) L(\mathbf{p}, \omega_i) d\omega_i + \frac{P_{f_s}(\mathbf{p}) \sigma_s(\mathbf{p})}{\sigma_{\text{maj}}} A(\mathbf{p}) \int_{S^2} f_s(\omega_i, \mathbf{p}, \omega) L(\mathbf{p}, \omega_i) d\omega_i \quad (3)$$

$L_e(\mathbf{p}, \omega)$ represents the radiance emitted from point \mathbf{p} in the direction ω . All coefficients in Equ. 3 can be multichannel and used to describe spectrally varying media.

3.2. Solution Operator

During the recursive sampling process for VPT using 3, the calculation of illumination contributions is limited to cases where the camera ray intersects with a light source or luminous voxels. This approach was proved to be inefficient [14]. However, by employing an operator-based transformation [28,54,55], a more efficient estimation method can be obtained. The scattering operator can be defined as follows:

$$\begin{aligned} (\mathbf{V}^y h)(\mathbf{p}, \omega) &= \sigma_s(\mathbf{p}) \cdot \int_{S^2} y(\omega', \mathbf{p}, \omega) \left[\int_0^z T_r(\mathbf{p} \leftrightarrow \mathbf{q}) h(\mathbf{q}, \omega') dt \right] d\omega' \\ (\mathbf{N}_s h)(\mathbf{p}, \omega) &= \sigma_n(\mathbf{p}) \int_0^z T_r(\mathbf{p} \leftrightarrow \mathbf{q}) h(\mathbf{q}, \omega') dt \end{aligned} \quad (4)$$

where $\mathbf{q} = \mathbf{p} + t\omega'$. Thus, the light transport can be described using radiance equilibrium [54,55]:

$$\begin{aligned} L_o &= \sigma_a(\mathbf{p}) L_e + (c^{f_s} \mathbf{V}^{f_s} + c^\rho \mathbf{V}^\rho + \mathbf{N}_s) L_o \implies L_o \left(I - (c^{f_s} \mathbf{V}^{f_s} + c^\rho \mathbf{V}^\rho + \mathbf{N}_s) \right) = \sigma_a(\mathbf{p}) L_e \\ \implies L_o &= \left(I - (c^{f_s} \mathbf{V}^{f_s} + c^\rho \mathbf{V}^\rho + \mathbf{N}_s) \right)^{-1} \sigma_a(\mathbf{p}) L_e \implies L_o = \sum_{k=0}^{\infty} (c^{f_s} \mathbf{V}^{f_s} + c^\rho \mathbf{V}^\rho + \mathbf{N}_s)^k \sigma_a(\mathbf{p}) L_e \end{aligned} \quad (5)$$

where I is an identity transformation, $c^{f_s} = P_{f_s}(\mathbf{p}) A(\mathbf{p})$ and $c^\rho = P_\rho(\mathbf{p})$.

Intuitively, the expression $(c^{f_s} \mathbf{V}^{f_s} + c^\rho \mathbf{V}^\rho + \mathbf{N}_s) L_e$ can be interpreted as the direct illumination radiance received at \mathbf{p} , and then scattering (including null scattering) along the direction ω . In the field of DVR, $\sigma_a(\mathbf{p}) L_e$ represents opacity-weighted emission [14]. Building upon Equ. 5, the estimation of $L(\mathbf{p}, \omega)$ can be decomposed into the sum of radiances along the direction ω reaching point \mathbf{p} after different numbers of scattering (including null scattering) events. This decomposition delineates the

light reaching the camera after scattering various numbers of times. VPT sampling represents the inverse process of light transport, where sampling rays accumulate contributions of direct illumination at the locations of each real scattering event.

3.3. Self-Occlusion

Shadow generation often encounters the phenomenon of “self-occlusion”, when using VPT in volumetric data containing implicit surface structures, as illustrated in Figure 2. For surface model rendering, light reflected from objects captured by the virtual camera includes both surface and subsurface scattering [14]. However, in DVR, sampling rays traverse implicit surfaces within the volumetric data, where the light received by sampling rays is obstructed by these implicit surfaces, resulting in darkened rendering outcomes.

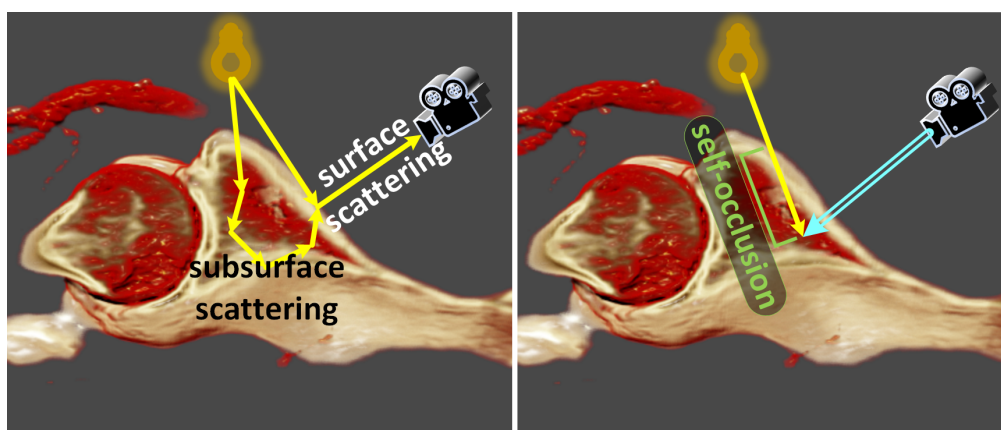


Figure 2. The left image depicts light scattered to the virtual camera, encompassing both surface and subsurface scattering components. However, in VPT, sampling rays may unavoidably penetrate directly into the interior of the object instead of precisely sampling the object’s implicit surface, as shown in the right image. Consequently, occluded shadow rays result in darkened renderings.

Although increasing the brightness of the light source can enhance the brightness of the rendering results, it compromises the realisticness of the results, as shown in Figure 3. Some iso-surface-based DVR methods, such as the DVR method used in [12], employ recursive techniques to precisely locate surface intersections before shading computation. However, this approach diminishes sampling efficiency and compromises the visual realisticness of the rendering results, particularly for subsurface scattering effects. We suggest multiplying the extinction coefficient σ_t of the entire volume by a factor $a \in (0.75, 0.85)$ when sampling shadow rays. Although this approach is somewhat heuristic, it enhances the material realisticness of the rendering results as shown in Figure 3.

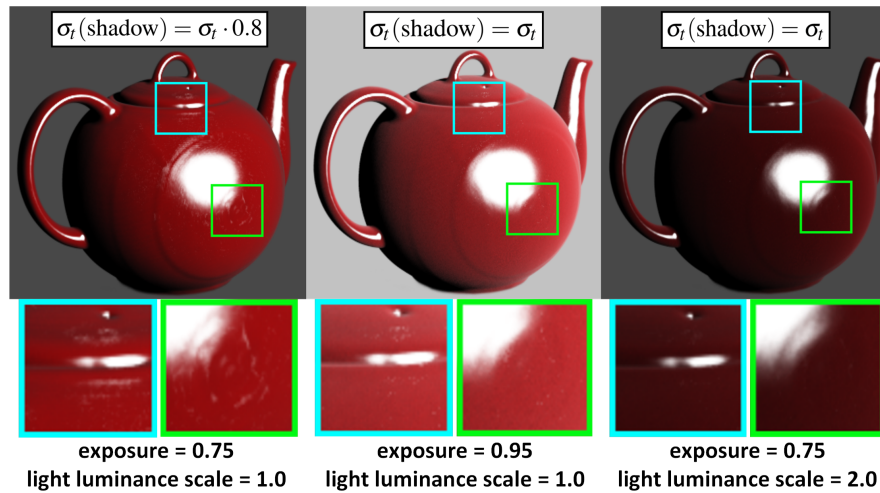


Figure 3. In these figures, $\sigma_t(\text{shadow})$ denotes the extinction coefficient used for sampling shadow rays. The left image depicts the rendering result obtained using our method, which involves multiplying σ_t by 0.8 to obtain $\sigma_t(\text{shadow})$. The middle and right images represent sampling of free paths and shadow rays using the same extinction coefficient. The middle image enhances the overall exposure, albeit at the expense of losing shading details. Despite doubling the brightness of the light source in the right image, the rendering output remains dim.

4. Hierarchical Volume Grid Accelerator

Sampling efficiency can be enhanced through the use of acceleration structures. Notable examples of these accelerators include the octree [32] and the Volume Dynamic B+Tree (VDB) [56]. Some general-purpose volumetric accelerators do not facilitate real-time structure updates on GPUs [56]. A common approach to accelerating DVR involves partitioning the volumetric space into multiple grids and recursively sampling each grid using the 3D-DDA algorithm [57,58]. However, this method does not support skipping large empty areas.

Xu et al. [53] proposed an acceleration structure that combines octrees and macro cells; however, this structure does not fully leverage the sampling efficiency of volumetric hierarchies. Additionally, the design for updating this acceleration structure is relatively simplistic, as it only considers one-dimensional transfer functions. In their approach, each node stores the minimum and maximum voxel values within it, and if the transfer function changes, the maximum extinction coefficient stored at the node is updated based on the voxel range defined by these min-max values. This method is inadequate for two reasons: first, the maximum extinction coefficient calculated from the voxel range may exceed the actual maximum extinction coefficient for the node, resulting in low sampling efficiency; second, it is not applicable for more complex or multidimensional transfer functions [10].

Our multi-hierarchical volumetric grid is illustrated in Figure 4. The octree is stored as a list, with each node's information including the spatial boundary range, parent node index, indices of the eight child nodes, and the maximum extinction coefficient within that node. The leaf nodes store the index range of the volume data.




Figure 4. The octree is constructed in a top-down manner, starting from Level 0 and progressively building downward, with the octree information organized in a list format. Conversely, when the transfer function changes, the octree is updated in a bottom-up fashion: leaf nodes are updated first, followed by subsequent updates moving upward.

When the transfer function changes, the acceleration structure should be updated. To achieve more accurate updates while accommodating various forms of transfer functions, a new update process is designed, as illustrated in Figure 4. First, leaf nodes are updated in parallel using CUDA, with each leaf node assigned a separate CUDA thread. Each thread traverses all voxels within the leaf node and updates the maximum extinction coefficient based on the extinction coefficient calculated from the transfer function. For levels topper than level 4, CPU parallel updating is employed, while higher levels continue to use CUDA. Nodes from Level 5 nodes is copied to CPU memory for updating the toppe-level hierarchy. This update strategy enables the maximum level 7 volumetric octree to be updated in just 15 ms for 512^3 resolution data, meeting real-time requirements. Figure 5 shows an example of the multi-hierarchical volumetric grid. The first row shows the volumetric spatial range of each node at different levels, while the second row visualizes the nodes containing valid voxels (those with opacity greater than 0) at each level.




Figure 5. The multi-hierarchical volumetric grid. The volumetric grid is organized in the form of an octree. The first row shows the volumetric spatial range occupied by each node at every level. The second row visually distinguishes the nodes containing valid voxels at each level by displaying them in different colors.

To enhance sampling efficiency, we developed a new multi-hierarchical volumetric sampling method. The execution process for volumetric traversal is guided by the following principles:

- If the current node does not contain valid voxels, the query proceeds to the parent node. If the parent node also lacks valid voxels, the search continues upward. This process continues until a

node is found that lacks valid voxels while its parent node contains valid voxels, at which point the sampling ray will bypass the current node.

- If the current node contains valid voxels and the maximum attenuation coefficient is less than 0.2, the search will proceed upward. If the maximum attenuation coefficient exceeds 0.6, the search will move downward. When the maximum attenuation coefficient is between 0.2 and 0.6, or if the node is at level 0 or at the leaf level, 3D-DDA sampling will be performed. If the maximum attenuation coefficient of the current node is less than 0.2, but the maximum attenuation coefficient of the parent node exceeds 0.6, traversal will continue at the current node. The thresholds of 0.2 and 0.6 are derived from empirical observations.
- Since the volumetric grid size is uniform across all levels, the node index of any point in any level can be directly calculated. When a ray samples within a node, it is assumed that the next node belongs to the same level; however, the search direction—upward or downward—will be determined by the maximum attenuation coefficient of the current node.

The multi-hierarchical querying mechanism we designed significantly enhances the efficiency of volume traversal. First, during 3D-DDA traversal, calculating the ray's advance range within the current node can be relatively time-consuming. To address this, we implement a fast bounding box intersection method as described in [14], which reduces the need for extensive division calculations while still requiring a considerable amount of multiplication and addition. In contrast, the computational load associated with querying the parent node and locating child nodes is minimal. When a large area exhibits low density, 3D-DDA can be executed directly at higher-level nodes, thereby avoiding multiple intersection calculations with the lower-level nodes. Additionally, the maximum attenuation coefficient of a node may greatly exceed that of most of its child nodes, and the downward querying mechanism improves sampling efficiency in such cases. However, it is crucial to ensure that the level of the leaf nodes is neither too low nor too high, as this can highly impact sampling efficiency. Empirically, setting the leaf node level to 6 proves to be a reasonable choice.

5. Multi-Material Transfer Functions

The transfer function is utilized to map voxel values to material properties [10]. Simple transfer functions consider only low-dimensional features of the data, such as original voxel intensity values and gradient magnitudes. In contrast, complex transfer functions facilitate more targeted voxel classification based on manual user interaction. The forms of these transfer functions vary in their classification methods. We categorize all forms of transfer functions into the following five categories: (1) Coefficient curve transfer function (CCTF), (2) Coefficient piecewise constant transfer function (CPCTF), (3) Coefficient discrete point transfer function (CDPTF), (4) Material piecewise constant transfer function (MPCTF), (5) Material discrete point transfer function (MDPTF).

Each category can be applied to a proxy volume for voxel classification. A proxy volume is generated by pre-processing the original volume, which itself is a proxy volume. Common proxy volumes include gradient magnitude volumes, opacity volumes, and mask volumes generated by segmentation. However, when multiple proxy volumes are present, conflicts may arise between the transfer function and the proxy volume. Thus, a management tool is necessary to regulate these situations, as described later.

First, the proxy volume is described. Any low-dimensional or high-dimensional volumetric feature can be represented as a proxy volume, which exists in two forms: volumes with continuous voxel value changes, such as gradient magnitude volumes [31] and soft-labeled volumes [42] obtained through segmentation; and volumes with discrete voxel values, such as hard-labeled volumes derived from semantic segmentation [43]. Different classification methods should be applied to proxy volumes with continuous changes and those with discrete values. Specifically, volumes represented in discrete form should utilize nearest neighbor interpolation during sampling to obtain current voxel values, rather than linear interpolation, to avoid incorrect and undefined intermediate values [1,43].

The coefficient transfer function assumes that the material of the controlled proxy volume is unique, but material properties such as roughness and reflectance are variable. This type of transfer function solely adjusts the material coefficients [31]. The relationship between the coefficient values and the voxel values of the proxy volume can take three forms: curve, piecewise constant, and discrete. In contrast, the material transfer function does not assume the uniqueness of the material in the controlled proxy volume. Since material types are discrete quantities, without intermediate transitions like coefficient values, the material transfer function is limited to either piecewise constant or discrete forms. For instance, in [49], the transfer function's inputs include volumetric opacity values, assigning a glass material to areas of low opacity and a metallic material to areas of high opacity. This can be implemented by generating an opacity proxy volume and then applying an MPCTF to it. Another example of material transfer functions [12], which assign different discrete material types based on mask values, facilitate rendering by recursively identifying the boundaries between different materials. The effects of various transfer functions are illustrated in Figure 6 and Figure 7.




Figure 6. Visualization of cell volume generated via electron microscopy. The first two images from the left illustrate the classification effects on the original volume and the mask volume using a one-dimensional transfer function. The distinction between the third and fourth images lies in the different transfer functions applied to the mask volume. Notably, using a CDCTF for the discrete value mask yields a clearer delineation.




Figure 7. The visualization of brain tumors and tissue edema is presented. The left and middle images employ a two-dimensional transfer function, applying the same type of transfer function to both the original volume and the mask volume, resulting in the target material coefficients through their multiplication. In contrast, the right image utilizes different material types for different mask values and applies the CCTF to further classify the coefficients for each material type.

When multiple transfer function types are applied to a single volumetric dataset, their compatibility must be considered. For instance, in VTK [59], multi-dimensional transfer functions compute target coefficients by multiplying the coefficient calculated in each dimension. In cases of conflict between two types, a priority must be established. For example, in DVR of segmented results, if the original volume employs a CCTF and the mask volume uses a MDPTF, the classification method can be set to prioritize MDPTF when the opacity controlled by MDPTF exceeds zero; otherwise, CCTF is used as the classification method. As shown in the right image in Figure 7, classification using the material transfer function and coefficient transfer function simultaneously yields a visualization with significantly greater detail compared to simple coefficient multiplication.

In summary, our transfer function technique is defined as comprising three steps: first, the generation of proxy volumes; second, the application of five types of transfer functions to assign material properties to these proxy volumes; and third, the merging of the material properties obtained from each proxy. Various complex transfer function techniques in [10], including the use of semantic segmentation for pre-processing volumetric data, can be categorized under the first step—generating proxy volumes. This categorization allows for the compatibility of diverse complex transfer function techniques within a single DVR framework.

6. Implementation and Evaluations

In this section, we first present some implementation details of the renderer, followed by an evaluation of our techniques. While we have already included some comparisons of various effects in the Sec. 5, our evaluation encompasses three additional aspects: first, we validate the effectiveness of our system in surgical planning through two medical applications. Our system is compared with several state-of-the-art (SotA) realistic volume rendering techniques; second, we conduct a user study to assess the realism of renderings produced by different photorealistic volumetric rendering algorithms; third, the sampling efficiency of our acceleration structure is compared with that of conventional acceleration structures; fourth, we tested the impact of different channel numbers of extinction coefficient on the rendering results.

6.1. Implementation Details

We implemented a physically-based volume renderer, with the rendering pipeline executed on CUDA. Compared to existing open-source photorealistic volume renderer [31], our approach features several key improvements: (1) The ExposureRenderer is based on ray marching sampling techniques. We developed a VPT-based renderer using the MM-RTM proposed in this paper, which supports both single-channel and multi-channel extinction coefficients. (2) Due to the limited interpolation capabilities of CUDA for 16-bit integer textures, we opted to use 32-bit floating-point textures instead. Our renderer supports the simultaneous loading of four proxy volumes. That is, the volume data is generated as *float4*-type textures. (3) The multi-hierarchical volumetric acceleration structure proposed in this paper is implemented.

Experimental hardware. All experiments were conducted on a PC workstation equipped with an Intel i7-12700KF CPU, 32 GB of RAM, and an NVIDIA GeForce RTX 4070 Ti GPU boasting 12 GB of video memory.

6.2. Comparison with the SotA in Surgical Planning

The most important application scenario for realistic DVR is medical image visualization, where the rendered results need to be as beneficial as possible for medical use. In this section, the advantages of our system over SotA methods in preoperative surgical planning are demonstrated. While numerous approximation and acceleration techniques have been recently developed for DVR [40,41,50,51], our comparison primarily focuses on the rendering technology. Technologies that are not compared are mainly designed for real-time acceleration, rather than enhancing visualization quality; thus, they are not directly compared. The proposed system also implements schemes for real-time browsing: supplementary materials include scenes reconstructed using neural rendering techniques, demonstrating real-time browsing effects achieved with NeRF and 3DGS technologies.

The SotA researches utilized for comparison are outlined as follows. We endeavor to replicate the rendering results presented in these papers to the best we can do: (1) *VPL-DVR* [60] approximates global illumination by computing the direct illumination and single scattering from a set of virtual light sources. It efficiently handles transfer function and volume density updates by recomputing the contribution of these virtual lights and progressively updating their volumetric shadow maps and locations. (2) *RMSS* [22] (Ray-Marching-based Single-Scattering) is an optimized physically-based DVR technique developed based on the open-source ExposureRenderer [31] to enhance the efficiency of environment light sampling. ExposureRenderer employs the ray-marching technique for free path

sampling in volume space during its implementation. (3) AR^2T [49] represents a physically-based DVR technique supporting multiple scattering. In AR^2T , four material types are defined, and the material is determined based on the volumetric density at the current position. Dielectric material treatment is applied when the density falls below a certain threshold t_0 , while metallic material treatment is applied when the density exceeds a certain threshold t_1 .

This study explores the advantages of our realistic DVR framework in surgical planning, including two application scenarios: (1) CT imaging plays a crucial role in both the diagnosis and treatment of fractures. It provides detailed multi-layered and multi-angular images, helping physicians to accurately identify the location, fracture patterns, and fragment conditions. Three-dimensional visualization of the fracture site is an essential component of surgical planning. (2) Coronary artery stenosis, often caused by atherosclerosis, can lead to myocardial ischemia and even myocardial infarction. During interventional surgery, physicians first use a balloon to dilate the artery and then implant a stent to maintain vessel patency, ensuring adequate blood flow to the myocardium. Three-dimensional visualization of coronary angiography data is an indispensable part of surgical planning. The results obtained using different visualization methods for two distinct scenarios are shown in Figure 8.




Figure 8. The visualization of CT data scans for patients undergoing fracture repair surgery and coronary artery stent surgery is shown. The first and second rows display the preoperative and postoperative visualizations for the fracture surgery, respectively. The third row illustrates the preoperative visualization for the stent surgery, with green arrows indicating the site of arterial stenosis.

We invited 8 clinical medical experts to evaluate the visualization results of different methods in this user study. Three evaluation criteria were established to ensure fair assessments, as outlined in Table 2. After a ten-minute evaluation of rendering results, participants rated their experiences on a scale from 1 to 5 (a higher score indicates better performance). Our statistical analysis employed one-way analysis of variance (ANOVA) to test the null hypothesis of equal correctness means across techniques. A p -value of ≤ 0.05 indicated a statistically significant difference in the scoring results between different methods.

Table 2. The evaluation criteria for helping participants assess the applicability of different methods in surgical planning.

	Evaluation Criteria
<i>EC1</i>	the visualization results' ability to support surgical planning tasks.
<i>EC2</i>	the effectiveness of the visualization in presenting three-dimensional spatial information.
<i>EC3</i>	the ability of the visualization to help doctors quickly locate the lesion.
<i>EC4</i>	the willingness to adopt this technology in surgical planning.

The scoring statistical results of the user study are presented in Figure 9. The scoring results indicate that medical experts unanimously consider our method to be the most ideal visualization solution. They confirmed that our approach is sufficiently effective for application in surgical planning. The experts consistently rated our method highly and expressed a strong willingness to use it for future surgeries. It is important to note that, while realistic rendering technologies cannot achieve things that other traditional rendering techniques cannot, enhancing realism and rendering details significantly improves the comfort and spatial awareness of doctors during observation. Some experts noted that higher levels of realism in the visualization make them feel more at ease, find the visuals more aesthetically pleasing, and enable them to focus more on diagnosing and planning around the lesion. To sum up, the method proposed in this paper demonstrates superior applicability for surgical planning compared to existing methods.




Figure 9. Visualization of the user ratings evaluated for the four given criteria. The boxes and whiskers display the minimum, maximum and median (as a red horizontal line within the box).

The differences in realism among various visualization methods will be experimentally evaluated and discussed in the next section. It is important to note that the transfer function management system presented in this paper plays a crucial role in integrating various types of volume labels and enhanced information into the rendering results. In the case of coronary angiography visualization, a two-dimensional transformation function is applied, with the first dimension being the CCTF and the second dimension the MDPTF. The coronary arteries are segmented and assigned a highlighted material type using the MDPTF. To ensure fairness, all rendering techniques employed this transfer function approach. Furthermore, a detailed comparison of the differences in rendering techniques and their effects will be provided in Sec. 6.3.

6.3. Comparison of Realistic Effects

In the user study comparing the realism of different methods, 19 participants (8 females and 11 males) aged between 23 and 30 years were recruited. Among these individuals, 12 had taken courses in medical visualization and had some familiarity with realistic DVR. Before the user study, we provided a brief overview of realistic DVR techniques to the participants. The rendering results were presented to each participant in a randomly shuffled order to minimize bias. Participants were able to freely

switch between these four DVR methods and four volumetric datasets, as shown in Figure 10. High dynamic range (HDR) panorama environment maps were employed as light sources to enhance the material realisticness of the rendering results [11,14].




Figure 10. Visual comparison of realistic DVR results generated using various state-of-the-art techniques and our method.

Three evaluation criteria are as outlined in Table 3. Besides evaluating the realisticness of rendering results, we also assessed the integration of the rendering results with the environment. To ensure the presence of rating differentiation, we encourage users to distinguish their ratings for the results generated using different DVR methods. After a ten-minute evaluation of rendering results, participants rated their experiences on a scale from 1 to 5.

Table 3. The evaluation criteria for helping participants objectively and comprehensively evaluate different realistic DVR methods.

	Evaluation Criteria
EC1	The material appearance realisticness of rendering result.
EC2	Beneficial for the perception of spatial structure.
EC3	Integration with the background environment.

The scoring statistical results of the user study are presented in Figure 11. The ANOVA test revealed significant differences among the various realistic DVR methods (p -value < 0.001), as assessed by different criteria in Table 4. As shown in Figure 11, our realistic DVR results outperform those of other state-of-the-art DVR methods on all four datasets. Most users indicated that our rendering results significantly exhibit higher realisticness compared to other methods. To ensure consistency with the original rendering results of those referenced papers, we also asked users to refer to the rendered

images presented in the literature of these studies [22,49,60], and users consistently affirmed the higher material realisticness and expressiveness of rendering results generated using our method.




Figure 11. Visualization of the user ratings evaluated for the three given criteria. The boxes and whiskers display the minimum, maximum, first and median (as a red horizontal line within the box).

Table 4. The quantitative comparison of different realistic DVR methods. Apart from the F -value and p -value, each value represents the mean/standard deviation.

Method	EC1	EC2	EC3
VPL-DVR [60]	2.25/0.80	2.87/0.69	2.75/0.76
RMSS [22]	3.12/0.83	3.87/0.81	3.24/0.82
AR ² T [49]	2.32/1.17	2.97/1.06	2.44/1.01
Ours	4.62/0.48	4.78/0.49	4.46/0.49
F -value	174.25	168.53	139.48
p -value	<0.001	<0.001	<0.001

The VPL-DVR [60] stands out as a significant technical solution based on approximate DVR global illumination. Numerous approximate global illumination approaches used in DVR can be found in [16]. Approximation methods typically involve a trade-off between rendering quality and interactive speed. The ExposureRenderer by Kroes et al. [31] can be considered the forerunner of physically-based realistic DVR research, exerting an important inspiration on subsequent studies in realistic DVR. The RMSS [22] primarily focuses on optimizing environment light sampling efficiency, sharing a similar DVR framework with the ExposureRenderer. The AR²T [49], a comprehensive realistic DVR solution supporting multiple scattering, exhibits limitations in its simplistic approach to material models, thereby restricting the portrayal of complex material effects.

Instead of simply combining existing techniques, we constructed a comprehensive MM-RTM. Our model offers multiple support for both surface scattering and volumetric scattering, resulting in more realistic material effects. Furthermore, we analyzed common issues in VPT-based realistic DVR, such as self-occlusion and low rendering detail, and provided corresponding solutions to address these problems. Our rendering results exhibit superior material realisticness, as shown in Figure 10. For instance, the metallic texture in **Engine**, the subsurface scattering effect in tooth and bone, and the light and shadow details on the rendered results. High material realisticness is also beneficial for users' comprehension and spatial perception of volume data.

6.4. Comparison of Accelerators

In this section, we evaluate the construction speed of different volumetric acceleration structures and their sampling efficiency. In addition to enhancing sampling performance, we introduce innovations that enable real-time updates of our acceleration structure, even for DVR with complex transfer functions. We compare various volumetric acceleration methods, including macro-cell acceleration, the Residency Octree [61] and our proposed multi-hierarchical acceleration structure. Additionally, since our accelerator allows for the configuration of different node levels, the update and sampling speeds of leaf nodes at various levels are analyzed.

Some advanced volumetric accelerators, such as SparseLeap [62], can effectively skip empty space but struggle to compute compact upper bounds for extinction coefficients within each effective sampling distance. This technology is primarily applied to ray-marching-based DVR algorithms and is less suited for VPT, so we do not include it in our comparisons.

Comparisons were conducted from three aspects: (1) the time required to update the accelerator structure; (2) the time consumed to locate the nearest and farthest valid voxel positions along a ray from the camera; and (3) the time taken for global illumination estimation at 10 samples per pixel (spp). The two volumetric datasets we selected for testing are **spathorhynchus** (resolution $1024 \times 1024 \times 750$) and **stag_beetle** (resolution $832 \times 832 \times 494$) [33]. The test results are presented in Table 5.

Table 5. The time consumed with various configurations of accelerators while performing three tasks, measured in milliseconds (ms).

spathorhynchus dataset			
	<i>Update</i>	<i>Locate Nearest</i>	<i>GI</i>
<i>No</i>	0	8.23	121.71
<i>Macrocell</i>	7.42	7.65	84.51
<i>Residency</i>	135.53	5.37	64.82
<i>Ours-6</i>	11.14	5.62	60.63
<i>Ours-7</i>	13.68	6.59	72.90

stag_beetle dataset			
	<i>Update</i>	<i>Locate Nearest</i>	<i>GI</i>
<i>No</i>	0	8.15	105.57
<i>Macrocell</i>	7.16	6.94	78.73
<i>Residency</i>	122.63	5.17	58.15
<i>Ours-6</i>	10.31	5.22	59.26
<i>Ours-7</i>	12.45	6.12	68.14

As shown in Table 5, our acceleration structure demonstrates a significant advantage in real-time interaction with transfer functions. In terms of sampling efficiency, the Residency Octree [61] performs similarly to our method, which also employs a multi-hierarchical structure. However, the Residency Octree is designed for out-of-core rendering of large datasets and does not support real-time updates. In contrast, we store the octree as a sequential list, simplifying design and implementation while enabling real-time updates. Our method, being hierarchical, also supports out-of-core rendering, as illustrated in Figure 6.

6.5. Configurations of Different Channel Numbers

When simulating light transport through colored participating media, multi-channel σ_t is commonly used [14,45]. However, employing multi-channel albedo with single-channel σ_t can also generate a colored rendering effect. We will demonstrate that employing a single-channel σ_t can yield rendering results indistinguishable from those obtained with multi-channel σ_t . Furthermore, the use of a single-channel σ_t can present less noise than using multi-channel σ_t during the interaction.

In Figure 12, it can be observed that by adjusting the transfer function, employing a single-channel configuration yields results that are difficult to discern from those of the multi-channel configuration by naked eyes. The difference maps in Figure 12 are used to show the subtle differences between the rendering results of these two configurations.

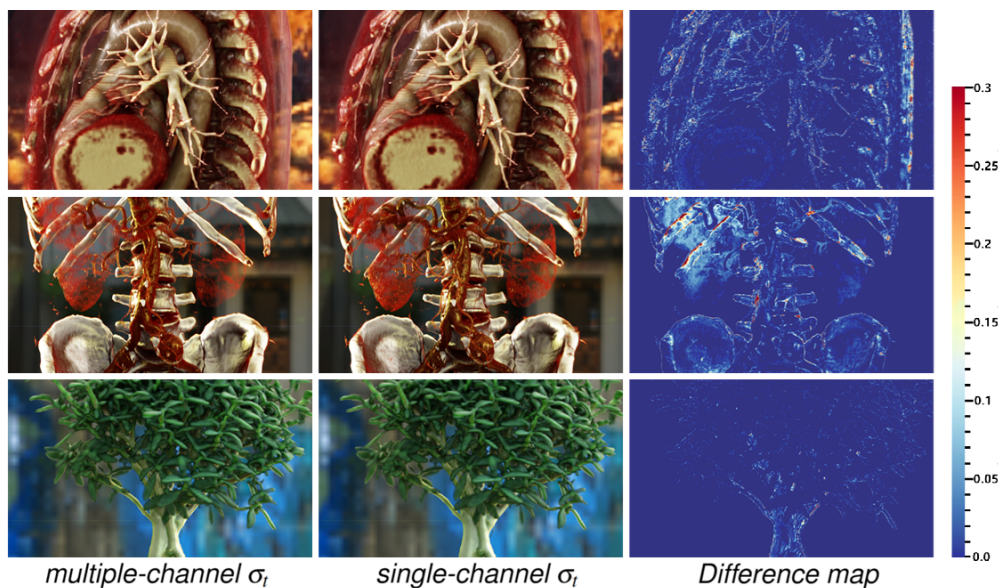


Figure 12. The rendering results are compared across different extinction coefficient σ_t channel configurations. The difference maps illustrate the acceptable absolute value discrepancies between the rendering results of single-channel configuration and multi-channel configuration.

The rendering noise levels across the two different channel configurations of σ_t were examined, as depicted in Figure 13. Single-channel configurations often exhibit lower noise levels when producing rendering effects similar to multiple-channel configurations. Therefore, we suggest that employing a single-channel σ_t with multi-channel albedo is a rational and effective choice in realistic DVR settings.

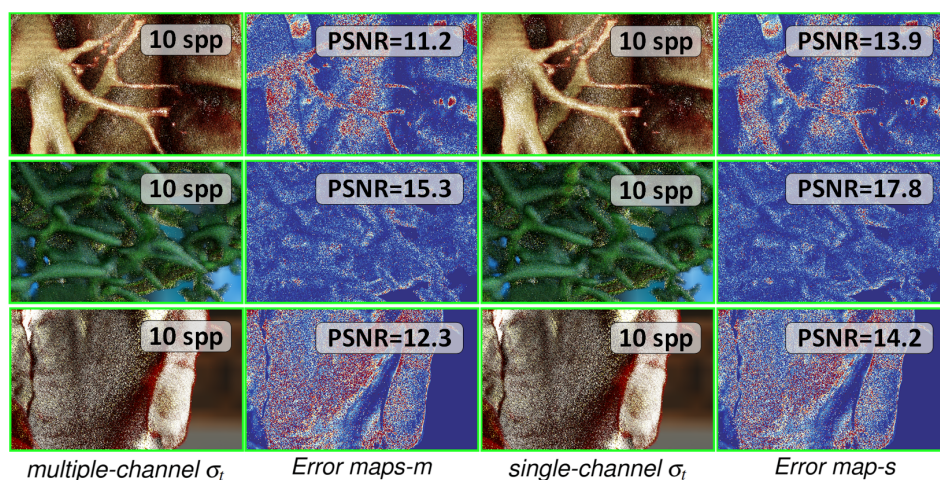


Figure 13. In the multi-channel version, *Error maps-m* denotes the discrepancy between rendering results with 10 spp and the reference results rendered with 4096 spp. *Error map-s* denotes the error in the single-channel rendition. A higher PSNR indicates lower image noise levels.

7. Conclusions and Future Work

In this paper, we present our realistic DVR technique with a novel multi-material radiative transfer model (MM-RTM). Building on advancements in light transport theories and simulation methods, we derive the theoretical framework for our MM-RTM, which efficiently incorporates multiple scattering effects. We also standardize various transfer function techniques and introduce five distinct forms of transfer functions along with proxy volumes, allowing our DVR framework to accommodate a variety of complex transfer function techniques effectively. Moreover, we develop a new multi-hierarchical volumetric acceleration method to improve sampling efficiency, which supports multi-level searches and volume traversal. Notably, our volumetric accelerator facilitates real-time structural updates during the application of complex transfer functions in DVR, enhancing the overall performance and

flexibility of our rendering framework. The effectiveness of our approach was evaluated via a user study. Our MM-RTM, the unified form of complex transfer functions, and the acceleration structure for real-time updates are complementary components that together form a comprehensive framework for realistic multi-material DVR. Our method was compared with existing state-of-the-art realistic DVR techniques. The result showed that our method generates the most realistic effects compared to these techniques.

During interactions, we discovered that achieving high realism also depends on the proper setup of light sources. Some HDR environment maps used as illumination sources may not result in a high level of material realism in DVR. Transfer function adjustments are also necessary to obtain the desired rendering appearance under different light sources. In future work, we will conduct in-depth research to explore the relationship between light source settings and the realism of rendering results. We will also try to innovate how the transfer function can be more conveniently and intuitively adjusted to obtain desirably optimized rendering results.

Author Contributions: Chunxiao Xu: Renderer Implementation, Conceptualization, Data curation, Methodology, Writing original draft. Xinran Xu: Conducting user studies, Experiments, Performing statistical analysis. Jiatian Zhang: Semantic segmentation of diseases and regions of interest, Performing data synthesis and transformation. Yiheng Cao: Revision of the manuscript. Lingxiao Zhao: Funding support & manuscript revision. All authors have read and agreed to the published version of the manuscript.", please turn to the [CRediT taxonomy](#) for the term explanation. Authorship must be limited to those who have contributed substantially to the work reported.

Funding: This work was supported in part by the National Natural Science Foundation of China (Nos. 82473472) and Suzhou Basic Research Pilot Project (SJC2021022).

Informed Consent Statement: Written informed consent has been obtained from the patients to publish this paper.

Data Availability Statement: The data used in this paper, except for those in Sec. 6.2, are all sourced from open datasets and . Data will be made available on request.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Preim, B.; Botha, C.P. *Visual computing for medicine: theory, algorithms, and applications*; Newnes, 2013. <https://doi.org/book/10.5555/2742301>.
2. Engel, K.; Hadwiger, M.; Kniss, J.M.; Lefohn, A.E.; Salama, C.R.; Weiskopf, D. Real-time volume graphics. In *ACM Siggraph 2004 Course Notes*; 2004; pp. 29–es.
3. Linsen, L.; Hagen, H.; Hamann, B. *Visualization in medicine and life sciences*; Springer, 2008. <https://doi.org/10.1007/978-3-540-72630-2>.
4. Dappa, E.; Higashigaito, K.; Fornaro, J.; Leschka, S.; Wildermuth, S.; Alkadhi, H. Cinematic rendering—an alternative to volume rendering for 3D computed tomography imaging. *Insights into imaging* **2016**, *7*, 849–856. <https://doi.org/10.1007/s13244-016-0518-1>.
5. Kruger, J.; Westermann, R. Acceleration techniques for GPU-based volume rendering. In *Proceedings of the IEEE Visualization, 2003. VIS 2003. IEEE, 2003*, pp. 287–292. <https://doi.org/10.1109/VISUAL.2003.1250384>.
6. Salama, C.R. Gpu-based monte-carlo volume raycasting. In *Proceedings of the 15th Pacific Conference on Computer Graphics and Applications (PG'07). IEEE, 2007*, pp. 411–414. <https://doi.org/10.1109/PG.2007.27>.
7. Hege, H.C.; Höllerer, T.; Stalling, D. Volume Rendering—Mathematical Models and Algorithmic Aspects. **1993**. <https://doi.org/10.20935/AcadEng7410>.
8. Max, N. Efficient light propagation for multiple anisotropic volume scattering. In *Proceedings of the Photorealistic Rendering Techniques*. Springer, 1995, pp. 87–104. https://doi.org/10.1007/978-3-642-87825-1_7.
9. Max, N. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics* **1995**, *1*, 99–108. <https://doi.org/10.1109/2945.468400>.

10. Ljung, P.; Krüger, J.; Groller, E.; Hadwiger, M.; Hansen, C.D.; Ynnerman, A. State of the art in transfer functions for direct volume rendering. In Proceedings of the Computer graphics forum. Wiley Online Library, 2016, Vol. 35, pp. 669–691. <https://doi.org/10.1111/cgf.12934>.
11. Zhang, Y.; Dong, Z.; Ma, K.L. Real-time volume rendering in dynamic lighting environments using precomputed photon mapping. *IEEE Transactions on Visualization and Computer Graphics* **2013**, *19*, 1317–1330. <https://doi.org/10.1109/TVCG.2013.17>.
12. Igouchkine, O.; Zhang, Y.; Ma, K.L. Multi-material volume rendering with a physically-based surface reflection model. *IEEE Transactions on Visualization and Computer Graphics* **2017**, *24*, 3147–3159. <https://doi.org/10.1109/TVCG.2017.2784830>.
13. Magnus, J.G.; Bruckner, S. Interactive dynamic volume illumination with refraction and caustics. *IEEE transactions on visualization and computer graphics* **2017**, *24*, 984–993. <https://doi.org/10.1109/TVCG.2017.2744438>.
14. Pharr, M.; Jakob, W.; Humphreys, G. Physically based rendering: From theory to implementation. **2023**.
15. Lindemann, F.; Ropinski, T. About the influence of illumination models on image comprehension in direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics* **2011**, *17*, 1922–1931. <https://doi.org/10.1109/TVCG.2011.161>.
16. Jönsson, D.; Sundén, E.; Ynnerman, A.; Ropinski, T. A survey of volumetric illumination techniques for interactive volume rendering. In Proceedings of the Computer Graphics Forum. Wiley Online Library, 2014, Vol. 33, pp. 27–51. <https://doi.org/10.1111/cgf.12252>.
17. Schott, M.; Pegoraro, V.; Hansen, C.; Boulanger, K.; Bouatouch, K. A directional occlusion shading model for interactive direct volume rendering. In Proceedings of the Computer Graphics Forum. Wiley Online Library, 2009, Vol. 28, pp. 855–862. <https://doi.org/10.1111/j.1467-8659.2009.01464.x>.
18. Sundén, E.; Ynnerman, A.; Ropinski, T. Image plane sweep volume illumination. *IEEE Transactions on Visualization and Computer Graphics* **2011**, *17*, 2125–2134. <https://doi.org/10.1109/TVCG.2011.211>.
19. Hadwiger, M.; Kratz, A.; Sigg, C.; Bühler, K. GPU-accelerated deep shadow maps for direct volume rendering. In Proceedings of the Graphics hardware, 2006, Vol. 6, pp. 49–52. <https://doi.org/10.1145/1283900.1283908>.
20. Kroes, T.; Post, F.; Botha, C. Interactive direct volume rendering with physically-based lighting. *Eurographics vi* **2012**, pp. 1–10.
21. Engel, K. Real-time Monte-Carlo path tracing of medical volume data. In Proceedings of the GPU technology conference, April, 2016, pp. 4–7.
22. von Radziewsky, P.; Kroes, T.; Eisemann, M.; Eisemann, E. Efficient stochastic rendering of static and animated volumes using visibility sweeps. *IEEE Transactions on Visualization and Computer Graphics* **2016**, *23*, 2069–2081. <https://doi.org/10.1109/TVCG.2016.2606498>.
23. Rowe, S.P.; Fishman, E.K. Fetal and placental anatomy visualized with cinematic rendering from volumetric CT data. *Radiology case reports* **2018**, *13*, 281–283. <https://doi.org/10.1016/j.radcr.2017.09.028>.
24. Berger, F.; Ebert, L.C.; Kubik-Huch, R.A.; Eid, K.; Thali, M.J.; Niemann, T. Application of cinematic rendering in clinical routine CT examination of ankle sprains. *American Journal of Roentgenology* **2018**, *211*, 887–890. <https://doi.org/10.2214/AJR.17.18630>.
25. Pattamapaspong, N.; Kanthawang, T.; Singsuwan, P.; Sansiri, W.; Prasitwattanaseree, S.; Mahakkanukrauh, P. Efficacy of three-dimensional cinematic rendering computed tomography images in visualizing features related to age estimation in pelvic bones. *Forensic science international* **2019**, *294*, 48–56. <https://doi.org/10.1016/j.forsciint.2018.10.003>.
26. Novák, J.; Selle, A.; Jarosz, W. Residual ratio tracking for estimating attenuation in participating media. *ACM Trans. Graph.* **2014**, *33*, 179–1. <https://doi.org/10.1145/2661229.2661292>.
27. Kutz, P.; Habel, R.; Li, Y.K.; Novák, J. Spectral and decomposition tracking for rendering heterogeneous volumes. *ACM Transactions on Graphics (TOG)* **2017**, *36*, 1–16. <https://doi.org/10.1145/3072959.3073665>.
28. Miller, B.; Georgiev, I.; Jarosz, W. A null-scattering path integral formulation of light transport. *ACM Transactions on Graphics (TOG)* **2019**, *38*, 1–13. <https://doi.org/10.1145/3306346.3323025>.
29. Misso, Z.; Li, Y.K.; Burley, B.; Teece, D.; Jarosz, W. Progressive null-tracking for volumetric rendering. In Proceedings of the ACM SIGGRAPH 2023 Conference Proceedings, 2023, pp. 1–10. <https://doi.org/10.1145/3588432.3591557>.
30. Galtier, M.; Blanco, S.; Caliot, C.; Coustet, C.; Dauchet, J.; El Hafi, M.; Eymet, V.; Fournier, R.; Gautrais, J.; Khuong, A.; et al. Integral formulation of null-collision Monte Carlo algorithms. *Journal of Quantitative Spectroscopy and Radiative Transfer* **2013**, *125*, 57–68. <https://doi.org/10.1016/j.jqsrt.2013.04.001>.

31. Kroes, T.; Post, F.H.; Botha, C.P. Exposure render: An interactive photo-realistic volume rendering framework. *PloS one* **2012**, *7*, e38586. <https://doi.org/10.1371/journal.pone.0038586>.
32. Knoll, A. A Survey of Octree Volume Rendering Methods. *VLUDS* **2006**, pp. 87–96.
33. Klacansky, P. *Open scientific visualization datasets*; 2013.
34. Rosset, A.; Ratib, O. *OsiriX DICOM Image Library*; 2002.
35. Müller, T.; Evans, A.; Schied, C.; Keller, A. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)* **2022**, *41*, 1–15. <https://doi.org/10.1145/3528223.3530127>.
36. Tancik, M.; Weber, E.; Ng, E.; Li, R.; Yi, B.; Wang, T.; Kristoffersen, A.; Austin, J.; Salahi, K.; Ahuja, A.; et al. Nerfstudio: A modular framework for neural radiance field development. In Proceedings of the ACM SIGGRAPH 2023 Conference Proceedings, 2023, pp. 1–12. <https://doi.org/10.1145/3588432.3591516>.
37. Kerbl, B.; Kopanas, G.; Leimkühler, T.; Drettakis, G. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* **2023**, *42*, 139–1. <https://doi.org/10.1145/3592433>.
38. Niedermayr, S.; Neuhauser, C.; Petkov, K.; Engel, K.; Westermann, R. Application of 3D Gaussian Splatting for Cinematic Anatomy on Consumer Class Devices. **2024**. <https://doi.org/10.2312/vmv.20241195>.
39. Ament, M.; Zirr, T.; Dachsbacher, C. Extinction-optimized volume illumination. *IEEE transactions on visualization and computer graphics* **2016**, *23*, 1767–1781. <https://doi.org/10.1109/TVCG.2016.2569080>.
40. Bauer, D.; Wu, Q.; Ma, K.L. Photon Field Networks for Dynamic Real-Time Volumetric Global Illumination. *IEEE Transactions on Visualization and Computer Graphics*. **2023**. <https://doi.org/10.1109/TVCG.2023.3327107>.
41. Wu, Q.; Bauer, D.; Doyle, M.J.; Ma, K.L. Interactive volume visualization via multi-resolution hash encoding based neural representation. *IEEE Transactions on Visualization and Computer Graphics*. **2023**. <https://doi.org/10.1109/TVCG.2023.3293121>.
42. Nguyen, N.; Bohak, C.; Engel, D.; Mindek, P.; Strnad, O.; Wonka, P.; Li, S.; Ropinski, T.; Viola, I. Finding nano-Ötzi: cryo-electron tomography visualization guided by learned segmentation. *IEEE Transactions on Visualization and Computer Graphics* **2022**, *29*, 4198–4214. <https://doi.org/10.1109/TVCG.2022.3186146>.
43. Engel, D.; Sick, L.; Ropinski, T. Leveraging Self-Supervised Vision Transformers for Segmentation-based Transfer Function Design. *IEEE Transactions on Visualization and Computer Graphics*. **2024**. <https://doi.org/10.1109/TVCG.2024.3401755>.
44. Li, M.; Jung, Y.; Song, S.; Kim, J. Attention-driven visual emphasis for medical volumetric image visualization. *The Visual Computer* **2024**, pp. 1–15. <https://doi.org/10.1007/s00371-024-03596-9>.
45. Novák, J.; Georgiev, I.; Hanika, J.; Jarosz, W. Monte Carlo methods for volumetric light transport simulation. In Proceedings of the Computer graphics forum. Wiley Online Library, 2018, Vol. 37, pp. 551–576. <https://doi.org/10.1111/cgf.13383>.
46. Fellner, F.A. Introducing cinematic rendering: a novel technique for post-processing medical imaging data. *Journal of Biomedical Science and Engineering* **2016**, *9*, 170–175. <https://doi.org/10.4236/jbise.2016.93013>.
47. 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* **2017**, *209*, 370–379. <https://doi.org/10.2214/AJR.17.17850>.
48. Rowe, S.P.; Johnson, P.T.; Fishman, E.K. Initial experience with cinematic rendering for chest cardiovascular imaging. *The British journal of radiology* **2018**, *91*, 20170558. <https://doi.org/10.1259/bjr.20170558>.
49. Denisova, E.; Manetti, L.; Bocchi, L.; Iadanza, E. AR2T: Advanced Realistic Rendering Technique for Biomedical Volumes. In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, 2023, pp. 347–357. https://doi.org/10.1007/978-3-031-43987-2_34.
50. Iglesias-Guitian, J.A.; Mane, P.; Moon, B. Real-time denoising of volumetric path tracing for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics* **2020**, *28*, 2734–2747. <https://doi.org/10.1109/TVCG.2020.3037680>.
51. Hofmann, N.; Martschinke, J.; Engel, K.; Stamminger, M. Neural denoising for path tracing of medical volumetric data. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* **2020**, *3*, 1–18. <https://doi.org/10.1145/3406181>.
52. Taibo, J.; Iglesias-Guitian, J.A. Immersive 3D Medical Visualization in Virtual Reality using Stereoscopic Volumetric Path Tracing. In Proceedings of the 2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR). IEEE, 2024, pp. 1044–1053. <https://doi.org/10.1109/VR58804.2024.00123>.
53. Xu, C.; Cheng, H.; Chen, Z.; Wang, J.; Chen, Y.; Zhao, L. Real-time Realistic Volume Rendering of Consistently High Quality with Dynamic Illumination. *IEEE Transactions on Visualization and Computer Graphics* **2024**, pp. 1–15. <https://doi.org/10.1109/TVCG.2024.3445339>.

54. Eric, V. Robust monte carlo methods for light transport simulation. *PhD thesis*. 1998.
55. Jakob, W.A. Light transport on path-space manifolds. *PhD thesis* 2013.
56. Museth, Ken. NanoVDB: A GPU-friendly and portable VDB data structure for real-time rendering and simulation. In Proceedings of the ACM SIGGRAPH 2021 Talks, 2021, pp. 1–2. <https://doi.org/10.1145/3450623.3464653>.
57. Amanatides, J.; Woo, A.; et al. A fast voxel traversal algorithm for ray tracing. In Proceedings of the Eurographics. Citeseer, 1987, Vol. 87, pp. 3–10. <https://doi.org/10.2312/egtp.19871000>.
58. Szirmay-Kalos, L.; Tóth, B.; Magdics, M. Free path sampling in high resolution inhomogeneous participating media. In Proceedings of the Computer Graphics Forum. Wiley Online Library, 2011, Vol. 30, pp. 85–97. <https://doi.org/10.1111/j.1467-8659.2010.01831.x>.
59. Hanwell, M.D.; Martin, K.M.; Chaudhary, A.; Avila, L.S. The Visualization Toolkit (VTK): Rewriting the rendering code for modern graphics cards. *SoftwareX* 2015, 1, 9–12. <https://doi.org/10.1016/j.softx.2015.04.001>.
60. Weber, C.; Kaplanyan, A.; Stamminger, M.; Dachsbacher, C. Interactive Direct Volume Rendering with Many-light Methods and Transmittance Caching. In Proceedings of the VMV, 2013, pp. 195–202. <https://doi.org/10.2312/PE.VMV.VMV13.195-202>.
61. Herzberger, L.; Hadwiger, M.; Krüger, R.; Sorger, P.; Pfister, H.; Gröller, E.; Beyer, J. Residency Octree: A Hybrid Approach for Scalable Web-Based Multi-Volume Rendering. *IEEE Transactions on Visualization and Computer Graphics* 2024, 30, 1380–1390. <https://doi.org/10.1109/TVCG.2023.3327193>.
62. Hadwiger, M.; Al-Awami, A.K.; Beyer, J.; Agus, M.; Pfister, H. SparseLeap: Efficient Empty Space Skipping for Large-Scale Volume Rendering. *IEEE Transactions on Visualization and Computer Graphics* 2018, 24, 974–983. <https://doi.org/10.1109/TVCG.2017.2744238>.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.