mathematics

MDPI

Article

A Voxel-Based Adaptive Exposure Model for Physically Based Volume Rendering

Yuliang Yuan, Jinzhu Yang *, Qi Sun and Yan Huang

School of Computer Science and Engineering, Northeastern University, Shenyang 110169, China; 2010668@stu.neu.edu.cn (Y.Y.); sampras418@hotmail.com (Q.S.); 8409huangyan@163.com (Y.H.)
* Correspondence: yangjinzhu@cse.neu.edu.cn

Abstract

In physically based volume rendering of volumetric medical imaging data, exposure levels are typically set empirically. However, differences in voxel values and opacity across transfer functions can make a fixed setting produce underexposed or overexposed results. To address this issue, we propose a voxel-contribution-based adaptive exposure model for physically based volume rendering. During volume integration, the method visibility-weights effective volumetric-sample energies using accumulated transmittance and opacity and estimates a global exposure coefficient from their weighted geometric mean. Compared with a two-dimensional HDR method operating on the rendered image, the proposed method incorporates three-dimensional information from volume integration into exposure estimation. The mean exposure-coefficient deviation is approximately 0.005%, compared with 0.031% for the two-dimensional method, while output mean-luminance variation remains below 1.10% over the tested range of normal illumination. The results demonstrate stable exposure and structural response across the evaluated transfer functions and illumination conditions.

Keywords: physically based volume rendering; adaptive exposure; volumetric rendering; radiative energy propagation; HDR

MSC: 68U05

1. Introduction

CT, MR, and PET examinations produce volumetric data that are commonly explored with three-dimensional visualization. Established volume-rendering methods include ray casting [1], splatting [2], and shear–warp rendering [3]. Interactive implementations often use local illumination models such as Phong [4], Blinn–Phong, or Lambertian shading; these models are efficient but do not simulate full light transport.

Ray tracing, path tracing [5], and photon mapping [6] provide more complete light-transport models. In volume rendering, however, the same lighting can produce markedly different brightness when the transfer function changes because opacity accumulation and transported energy also change. A fixed exposure may therefore be suitable for one transfer function and poor for another.

Adaptive exposure [7] and HDR tone reproduction [8] usually estimate brightness from a completed two-dimensional image. They can correct the displayed image, but the statistic no longer contains the individual transmittance, opacity, depth, and energy terms available during volume integration.

Check for updates

Academic Editors: Xiaoqiang Zhu and Lihua You

Received: 3 August 2026
Revised: 15 October 2026
Accepted: 27 October 2026
Published: 4 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

Increasing samples per pixel (SPP) reduces Monte Carlo noise, but does not by itself regulate the global exposure associated with a transfer function. Exposure estimation and sampling convergence are therefore evaluated separately in this work.

We estimate exposure from voxel contributions collected before projection. "Voxel-based" describes the source of the statistic; the method produces one global coefficient, not a per-voxel exposure field. The main contributions are:

- We formulate global exposure as energy normalization over pre-projection volumetric-sample contributions.
- We estimate the scale with a visibility-weighted geometric mean, $w_i = T_i\alpha_i$, and apply the resulting coefficient once to all RGB contributions. This preserves non-zero per-channel contribution ratios before clipping and quantization.
- We evaluate the method across SPP levels, transfer functions, illumination levels, and CT, MR, and PT volumes. The study includes image-space baselines, parameter tests, an averaging-operator ablation, and cost analysis. The mean exposure-coefficient deviation is 0.005%, compared with 0.031% for SGM, and luminance variation remains below 1.10% over the tested illumination range.

2. Related Work

Most physically based volume-rendering research focuses on transport, sampling, and illumination. Exposure is commonly set manually or estimated after projection, leaving comparatively little work on statistics collected during volume integration.

Early physically based approaches, including the Monte Carlo volume-rendering work of Csebfalvi et al. [9,10] and the photon-mapping method of Zhang et al. [11], improved volumetric lighting and shadows. Their focus was transport and rendering quality rather than automatic exposure estimation from volumetric samples.

More recent work, such as the path-traced framework of Iglesias-Guitian et al. [12], improves visual fidelity at low sample counts but does not formulate a volumetric exposure statistic.

Image-space tone reproduction includes global operators, such as Reinhard's photographic method, and local operators, such as the bilateral-filter method of Durand and Dorsey. Reinhard's method scales image luminance from its log-average and a middle-gray key before highlight compression [13]. Durand–Dorsey separates log luminance into base and detail layers, compresses the base, and recombines the result [14]. Both operate after projection and, therefore, have no access to the transmittance, opacity, depth, or sample contributions used during volume integration.

Dinesha et al. studied dynamic-range representation in HDR volume rendering [15], and Xu et al. addressed dynamic illumination under changing light conditions [16]. Neither work estimates global exposure from effective sample energies during volume integration.

These studies motivate a separate question: how should a global exposure scale be estimated from the information already available during volume integration?

3. Voxel-Based Adaptive Exposure Formulation

To address exposure inconsistencies in physically based volume rendering, we view exposure estimation as a volumetric energy normalization problem rather than a purely image-space intensity adjustment problem. The model estimates a single global exposure coefficient from the energies of participating volumetric samples during volumetric transport and applies this coefficient in volumetric accumulation.

This section describes the volumetric exposure model and its theoretical formulation.

Table 1 lists the notation. Here, a "voxel contribution" is evaluated at an effective ray-marching sample in the voxel grid.

Table 1. Notation used in the exposure formulation.

Symbol	Definition
$\mathbf{C}_j, \alpha_j$	RGB color and opacity evaluated at volume sample j.
T_j	Accumulated transmittance immediately before sample j.
E_0, E_j	Initial incident radiant energy and radiant energy arriving at sample j, respectively.
$\mathcal{E}_i$	Local volumetric radiative energy at effective sample i, before view-dependent visibility weighting.
$w_i = T_i\alpha_i$	View-dependent visibility weight of effective sample i.
n, K	Number of samples on one ray and total number of effective samples in the exposure statistic, respectively.
μ_g	Visibility-weighted geometric mean of local volumetric radiative energy.
G_a, ϵ	Target gray value and positive numerical-stability constant.
δ	Global exposure coefficient estimated in the volumetric domain.

3.1. Exposure Formulation Analysis

In physically based volume rendering, color accumulation along a viewing ray is commonly described by the emission–absorption model. Let $\mathbf{C}_j$ and α_j be the RGB color and opacity at volume sample j. Front-to-back accumulation gives

$$\mathbf{C}_{\text{image}}(x, y) = \mathbf{C}_0\alpha_0 + \sum_{j=1}^{n-1}\left[\left(\prod_{k=0}^{j-1}(1-\alpha_k)\right)\alpha_j\mathbf{C}_j\right] \tag{1}$$

where n denotes the number of contributing volumetric samples along the viewing direction.

The accumulated transmittance immediately before sample j is

$$T_j = \prod_{k=0}^{j-1}(1-\alpha_k), \qquad T_0 = 1. \tag{2}$$

The RGB contribution of sample j is therefore

$$\mathbf{C}_j^{\text{contrib}} = T_j\alpha_j\mathbf{C}_j \tag{3}$$

3.2. Voxel-Based Adaptive Exposure Model

The same front-to-back process gives a simple scalar energy relation. Let E_0 be the incident radiant energy and E_j the energy arriving at sample j. The energy arriving at the next sample is

$$E_{j+1} = E_j(1-\alpha_j), \qquad j = 0, \ldots, n-1. \tag{4}$$

The energy contribution at sample j is

$$E_j^{\text{contrib}} = E_j\alpha_j = E_0T_j\alpha_j \tag{5}$$

The absorbed contributions and the energy leaving the volume satisfy

$$\sum_{j=0}^{n-1} E_j^{\text{contrib}} + E_n = E_0\left[\sum_{j=0}^{n-1} T_j\alpha_j + \prod_{k=0}^{n-1}(1-\alpha_k)\right] = E_0. \tag{6}$$

According to the above formulation, all effective volumetric samples participating in integration can be identified through their non-zero transmittance and opacity contributions. Let K denote the total number of effective participating samples. To avoid applying the visible contribution in Equation (5) twice, $\mathcal{E}_i > 0$ in Equation (7) denotes the local volumetric

radiative energy of sample i before visibility weighting, and $w_i = T_i\alpha_i$; the corresponding visible contribution is $w_i\mathcal{E}_i$.

An unweighted geometric mean gives every effective sample equal influence. In semi-transparent data, many weakly visible samples may lower the mean and produce an excessively large exposure coefficient. The coefficient still scales all contributions uniformly; it does not selectively brighten deeper structures. We instead weight each sample by its visible contribution, $w_i = T_i\alpha_i$:

$$\mu_g = \exp\left(\frac{\sum_{i=1}^{K} w_i \log(\mathcal{E}_i + \epsilon)}{\sum_{i=1}^{K} w_i}\right) - \epsilon, \qquad \sum_{i=1}^{K} w_i > 0. \tag{7}$$

where ϵ is a small positive constant introduced to avoid numerical instability. The weighting is confined to the estimation of μ_g; the global exposure coefficient is subsequently defined by Equation (8) and uniformly scales the volumetric contributions.

Based on the geometric mean energy, a global exposure coefficient δ is introduced to regulate the overall brightness response:

$$\delta = \frac{G_a}{\mu_g} \tag{8}$$

where G_a denotes a target statistical operating point in the normalized linear RGB range. Photographic exposure commonly uses a middle-gray key near 0.18, motivated by surface reflectance and perceptual practice [13]; our statistic instead describes pre-projection volumetric radiative energy, so the meanings are not identical.

In volumetric rendering, however, energy is not determined solely by scene illumination. Opacity accumulation, transmittance attenuation, and transfer-function-dependent voxel responses all influence the final energy distribution. As a result, the statistical characteristics of volumetric energy differ from those of natural image luminance.

Exposure is estimated before projection from volumetric energy and visibility. We set $G_a = 0.5$, the midpoint of the normalized linear RGB interval, as the default operating point. It is the target of the volumetric statistic, not the arithmetic mean of the displayed pixels. Transfer functions, occlusion, color, clipping, and quantization still shape the final image. A fixed target may fail for extreme energy distributions; estimating G_a from the volumetric-energy distribution is left for future work.

The exposure-scaled RGB contribution of sample j is

$$\mathbf{f}_j = \delta T_j \alpha_j \mathbf{C}_j \tag{9}$$

The pixel color is then obtained by accumulating the scaled sample contributions:

$$\mathbf{P}_{\text{color}}(x, y) = \sum_{j=0}^{n-1} \mathbf{f}_j \tag{10}$$

Substituting the exposure-scaled contribution function into the accumulation equation yields

$$\mathbf{P}_{\text{color}}(x, y) = \delta \sum_{j=0}^{n-1} T_j \alpha_j \mathbf{C}_j \tag{11}$$

where n denotes the number of contributing samples along the viewing direction, as illustrated by the red line in Figure 1.

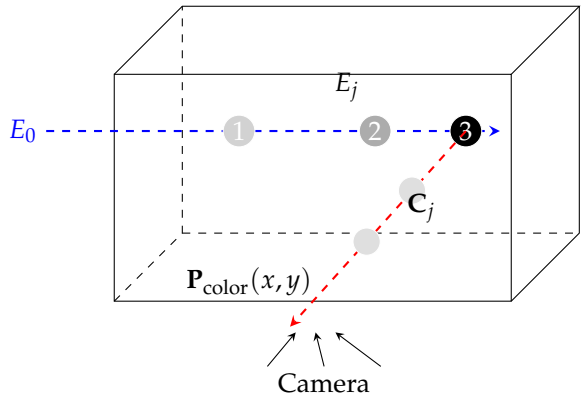


Figure 1. Volumetric color accumulation and energy transport diagram. The blue dashed arrow marks the direction of incident radiant-energy transport, where E_0 is the incident energy and E_j the energy arriving at sample j. The red dashed arrow marks the viewing direction toward the camera, along which the sample colors $\mathbf{C}_j$ are composited by opacity α_j and accumulated transmittance T_j and summed into the pixel color $\mathbf{P}_{\text{color}}(x, y)$.

Global Scaling.

Because δ uniformly rescales every sample contribution, it preserves their linear relationships before clipping and quantization. The operation is therefore a global normalization whose scale is estimated in the volumetric domain.

Stability Analysis.

Under a bounded volumetric-energy distribution, the exposure coefficient is derived from the geometric mean weighted by $w_i = T_i\alpha_i$, which incorporates both accumulated transmittance and the opacity of the current sample. Assume that the effective-sample set and its weights remain fixed and that only the local energies are perturbed. Let $\bar{w}_i = w_i / \sum_{k=1}^{K} w_k$. If only samples in a set S change from $\mathcal{E}_i$ to $\mathcal{E}'_i$, Equation (7) gives

$$\Delta \log(\mu_g + \epsilon) = \sum_{i \in S} \bar{w}_i \log\left(\frac{\mathcal{E}'_i + \epsilon}{\mathcal{E}_i + \epsilon}\right),$$

and therefore

$$\left|\Delta \log(\mu_g + \epsilon)\right| \leq \sum_{i \in S} \bar{w}_i \left|\log\left(\frac{\mathcal{E}'_i + \epsilon}{\mathcal{E}_i + \epsilon}\right)\right|.$$

Thus, small local energy fluctuations, or fluctuations confined to a region with a small accumulated visibility weight, have a correspondingly limited influence on the global exposure statistic. If the perturbation also changes α_i, T_i, or the effective-sample set, the weights change and this bound no longer applies directly; this case is evaluated in the transfer-function robustness experiment in Section 4.6.

The experiments use a multi-light configuration intended to show the overall anatomy under ordinary volume-rendering illumination. The method adjusts the display range of existing contributions; it does not create structural information. Extreme energy distributions are outside the tested operating range and are discussed as a limitation.

3.3. Exposure Consistency Analysis

The same positive coefficient is applied to every contributing sample. The following analysis concerns the linear floating-point result before clipping and quantization.

Let the original RGB contribution of sample j be

$$\mathbf{q}_j = T_j\alpha_j\mathbf{C}_j \tag{12}$$

After applying the proposed exposure coefficient, the normalized contribution becomes

$$\mathbf{q}'_j = \delta \mathbf{q}_j \tag{13}$$

For any two contributing samples j and k, and any color channel $c \in \{r, g, b\}$ with $q_{k,c} \neq 0$, the ratio satisfies

$$\frac{q'_{j,c}}{q'_{k,c}} = \frac{\delta q_{j,c}}{\delta q_{k,c}} = \frac{q_{j,c}}{q_{k,c}}. \tag{14}$$

Thus, the coefficient preserves each non-zero per-channel contribution ratio and rescales the accumulated RGB vector by a constant factor.

Similarly, the total accumulated contribution can be written as

$$\mathbf{Q} = \sum_{j=0}^{n-1} \mathbf{q}_j \tag{15}$$

and the corresponding exposure-normalized contribution becomes

$$\mathbf{Q}' = \sum_{j=0}^{n-1} \mathbf{q}'_j = \delta \sum_{j=0}^{n-1} \mathbf{q}_j = \delta \mathbf{Q}. \tag{16}$$

This is a global linear normalization, not a local modification of individual sample contributions. Any later clipping or quantization may still alter displayed ratios.

Furthermore, the formulation leaves volumetric occlusion unchanged. Since sample visibility is determined by accumulated transmittance,

$$T_j = \prod_{k=0}^{j-1} (1 - \alpha_k), \tag{17}$$

sample visibility is independent of the exposure coefficient. If the first sample is fully opaque,

$$\alpha_0 = 1, \tag{18}$$

the transmittance of all subsequent samples is

$$T_j = 0, \qquad j > 0. \tag{19}$$

All subsequent sample contributions then vanish for any exposure coefficient. Exposure normalization therefore cannot reveal a fully occluded structure.

Before display clipping and quantization, exposure normalization changes the overall level without changing non-zero per-channel contribution ratios or transmittance-based visibility. Visibility weighting affects only the estimate of the coefficient; the subsequent scaling is uniform.

4. Experimental Analysis

Experiments were run on an Intel Core i9-13900K workstation with 64 GB of RAM and an NVIDIA GeForce RTX 4090 GPU. All methods used the same empty-space skipping and spatial acceleration. Each comparison fixed the output resolution, camera, transfer function, clipping box, and SPP. Unless stated otherwise, $G_a = 0.5$, $\epsilon = 10^{-6}$, and the remaining-photon-energy cutoff was 1. The three lights used the same initial photon energy within each condition. Sampling was tested from 2 to 128 SPP.

The main experiments use one chest CT volume with three transfer functions: opaque bone, semi-transparent lung tissue, and a combined bone–lung display. They represent surface-dominated, low-opacity, and multi-tissue rendering conditions.

Three rectangular area lights are defined relative to the clipped-volume center $\mathbf{c}$, view direction $\mathbf{d}$, screen-right vector $\mathbf{r}$, screen-up vector $\mathbf{u}$, and maximum clipped-box extent L. Table 2 gives their positions, targets, areas, and colors. These settings remain fixed when photon energy is varied.

Table 2. Three-light configuration used in the experiments. Coordinates and sizes are normalized by the clipped-volume extent L.

Light	Center Position	Aim Target	Area $(w \times h)$	RGB Color
Key	$\mathbf{c} - 0.80L\mathbf{d} - 0.60L\mathbf{u} + 0.60L\mathbf{r}$	$\mathbf{c} - 0.15L\mathbf{u}$	$0.12L \times 0.12L$	$(1, 1, 1)$
Fill	$\mathbf{c} - 0.75L\mathbf{d} - 0.45L\mathbf{u} - 0.85L\mathbf{r}$	$\mathbf{c} - 0.05L\mathbf{u}$	$0.45L \times 0.45L$	$(0.12, 0.12, 0.12)$
Front fill	$\mathbf{c} - 0.90L\mathbf{d} - 0.25L\mathbf{u}$	$\mathbf{c} - 0.05L\mathbf{u}$	$0.80L \times 0.80L$	$(0.012, 0.012, 0.012)$

4.1. Experimental Data and Comparison Protocol

Sampling sensitivity is measured on the same $512 \times 512 \times 268$ lung CT scan at $\{2, 4, 8, 16, 32, 64, 128\}$ SPP. Adaptive exposure is enabled for all three transfer functions at every SPP. Camera, clipping range, lighting, photon energy, and display encoding remain fixed, and the 128-SPP image of each transfer function serves as its highest-sampled reference. Figure 2 shows the qualitative volume-rendering results and the adaptive-exposure outputs from 2 to 128 SPP, and Table 3 reports the corresponding benchmark relative to each 128-SPP reference.

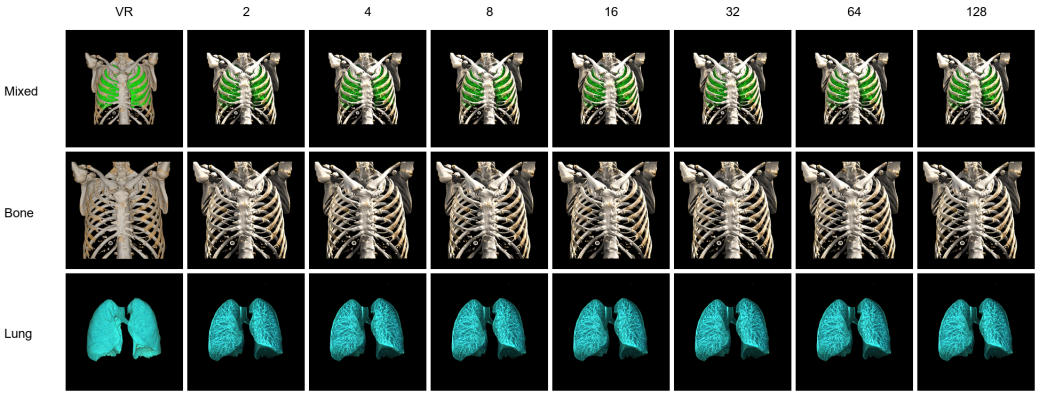


Figure 2. Qualitative VR results (first column) and adaptive-exposure results from 2 to 128 SPP for the mixed bone–lung, opaque-bone, and semi-transparent-lung transfer functions. The VR images are shown only for visual context and are not used in the numerical measurements; the 128-SPP result in each row serves as the highest-sampled reference for that transfer function.

Table 3. SPP benchmark for the three transfer functions relative to their respective 128-SPP results.

TF	SPP	PSNR (dB)	SSIM	Mean $\lvert\Delta L\rvert$	δ	Mean Luminance	Saturation (%)	Time (s)
Mixed	2	33.459	0.966371	3.136	1.698909	137.058	21.137	1.311
	4	35.266	0.976403	2.458	1.698456	137.210	21.190	2.356
	8	37.180	0.983549	1.871	1.698420	137.358	21.205	4.378
	16	39.153	0.989127	1.380	1.698487	137.447	21.220	8.712
	32	41.389	0.993379	0.960	1.698542	137.521	21.253	17.117
	64	44.624	0.996734	0.582	1.698562	137.568	21.276	33.877
	128	∞	1.000000	0.000	1.698547	137.598	21.290	67.662

Table 3. *Cont.*

TF	SPP	PSNR (dB)	SSIM	Mean $\lvert\Delta L\rvert$	δ	Mean Luminance	Saturation (%)	Time (s)
Bone	2	33.557	0.960033	3.016	1.423781	138.338	15.152	0.754
	4	35.024	0.970350	2.425	1.423752	138.489	15.147	1.313
	8	36.632	0.978317	1.913	1.423679	138.628	15.166	2.477
	16	38.635	0.984969	1.455	1.423619	138.821	15.206	4.562
	32	41.031	0.990466	1.043	1.423650	138.968	15.239	8.774
	64	44.141	0.995145	0.646	1.423629	139.051	15.247	17.410
	128	∞	1.000000	0.000	1.423645	139.090	15.233	34.625
Lung	2	36.198	0.977824	2.746	2.155638	105.793	8.223	1.438
	4	38.479	0.986030	2.104	2.155332	105.880	8.318	2.563
	8	40.933	0.991438	1.572	2.155416	105.972	8.424	5.022
	16	43.528	0.994828	1.133	2.155374	106.024	8.472	9.620
	32	46.290	0.997044	0.780	2.155287	106.071	8.495	19.045
	64	49.896	0.998568	0.473	2.155259	106.098	8.514	37.994
	128	∞	1.000000	0.000	2.155275	106.114	8.522	75.819

Error relative to 128 SPP decreases with SPP for all three transfer functions. At 64 SPP, each case exceeds 40 dB PSNR and 0.995 SSIM, with a mean absolute luminance difference below 1. We therefore use the 128-SPP images as the highest-sampled numerical references in the remaining experiments.

The main comparison uses two global estimators. Screen-Space Geometric Mean (SGM) computes the geometric mean luminance of nonzero HDR pixels after compositing and before exposure scaling, corresponding to the log-average exposure stage of photographic tone reproduction [13]. Nonlinear highlight compression and local contrast processing are omitted, so SGM produces only a uniform RGB scale. visibility-weighted geometric mean (VWGM) instead uses effective pre-projection samples weighted by $w_i = T_i\alpha_i$.

The evaluation covers SPP convergence, transfer-function and illumination robustness, the choice of averaging operator, parameter sensitivity, computational cost, and additional datasets. Durand–Dorsey tone mapping is included once as a local image-space reference, separate from the controlled SGM–VWGM comparison.

Metric definitions. *Exposure-coefficient deviation.* For a fixed method m and transfer function t, the relative deviation of the exposure coefficient at SPP s is measured against the coefficient obtained by the same method and transfer function at 128 SPP:

$$e_{\delta}^{m,t}(s) = 100\,\frac{|\delta_{m,t}(s) - \delta_{m,t}(128)|}{\delta_{m,t}(128)}.$$

For each method and transfer function, the reported summary is the arithmetic mean of these deviations over $s \in \{2, 4, 8, 16, 32, 64\}$; the 128-SPP result is used only as the reference. The numerical-stability experiment instead uses the result at $\epsilon = 10^{-6}$ under the same configuration as its reference.

Luminance CV. This metric measures the relative variation of display-domain output light intensity. Let $r_{p,c}, g_{p,c}, b_{p,c} \in [0, 1]$ be the normalized RGB channels after final display conversion and clipping, and let $F_c = \{p : \max(r_{p,c}, g_{p,c}, b_{p,c}) > 0\}$ be the nonblack-pixel set. Pixel output light intensity is the RGB mean multiplied by display range D; image mean output light intensity is its arithmetic mean over F_c:

$$L_{p,c} = D\,\frac{r_{p,c} + g_{p,c} + b_{p,c}}{3}, \qquad \bar{L}_c = \frac{1}{|F_c|}\sum_{p \in F_c} L_{p,c}.$$

The output-intensity CV is calculated across the compared conditions as

$$\mathrm{CV}_L = 100\,\frac{\sigma_{\mathrm{pop}}(\{\bar{L}_c\}_{c \in \mathcal{C}})}{\mathrm{mean}(\{\bar{L}_c\}_{c \in \mathcal{C}})}.$$

For the SPP experiment, $\mathcal{C} = \{2, 4, 8, 16, 32, 64, 128\}$; for the illumination experiment, $\mathcal{C} = \{0.8, 1.0, 1.3\}$. Multiplication by the constant display range D cancels between the numerator and denominator of the CV, so normalized $[0, 1]$ and 8-bit $[0, 255]$ representations give the same CV.

Saturation ratio. The saturation ratio is the percentage of nonblack foreground pixels for which at least one RGB channel reaches the upper display bound. Black background pixels are excluded, and differences between saturation ratios are reported in percentage points:

$$\mathrm{Sat}_c = 100\,\frac{|\{p \in F_c : \max(r_{p,c}, g_{p,c}, b_{p,c}) = 1\}|}{|F_c|}.$$

ROI-contrast CV. Display-domain luminance is first calculated as $Y_p = 0.2126R_p + 0.7152G_p + 0.0722B_p$. The same 160×160-pixel ROI is reused across all compared conditions of a given transfer function, and its black pixels are excluded. For each rendered image, normalized RMS contrast is

$$C_{\mathrm{ROI},c} = \frac{\sigma_{\mathrm{pop}}(Y_{R,c})}{\mathrm{mean}(Y_{R,c})}, \qquad \mathrm{CV}_{\mathrm{ROI}} = 100\,\frac{\sigma_{\mathrm{pop}}(\{C_{\mathrm{ROI},c}\}_{c\in\mathcal{C}})}{\mathrm{mean}(\{C_{\mathrm{ROI},c}\}_{c\in\mathcal{C}})}.$$

Thus, ROI-contrast CV measures the relative variation of ROI contrast across the corresponding experimental conditions. Mean ± SD summaries across the three transfer functions use the arithmetic mean and sample standard deviation.

4.2. SPP-Dependent Convergence Comparison

SGM and VWGM are evaluated at $\{2, 4, 8, 16, 32, 64, 128\}$ SPP for the mixed, bone, and lung transfer functions. PSNR, SSIM, mean absolute luminance difference, and relative coefficient error are computed against each method's own 128-SPP result. Figure 3 plots this convergence behavior, and Table 4 summarizes the complete sweep.

Table 4. Summary of the full SPP sweep for the primary comparison methods. Error sequences correspond to 2, 4, 8, 16, 32, and 64 SPP relative to each method's own 128-SPP result.

TF	Method	$\lvert\Delta\delta\rvert/\delta_{128}$ (%)	PSNR (dB)	SSIM	Luminance CV/Sat. (%)
		0.0628/0.0095	33.322/35.165	0.968333/0.977790	
		0.0430/0.0250	37.124/39.080	0.984700/0.989856	
Mixed	2D method	0.0004/0.0008	41.443/44.708	0.993852/0.997014	0.130/28.556
		0.0213/0.0054	33.459/35.266	0.966371/0.976403	
		0.0075/0.0035	37.180/39.153	0.983549/0.989127	
Mixed	Proposed	0.0003/0.0009	41.389/44.624	0.993379/0.996734	0.135/21.290
		0.0399/0.0597	33.647/35.145	0.961439/0.971312	
		0.0472/0.0324	36.751/38.778	0.979056/0.985431	
Bone	2D method	0.0278/0.0162	41.153/44.367	0.990724/0.995301	0.211/20.519
		0.0096/0.0075	33.557/35.024	0.960033/0.970350	
		0.0024/0.0018	36.632/38.635	0.978317/0.984969	
Bone	Proposed	0.0004/0.0011	41.031/44.141	0.990466/0.995145	0.194/15.233
		0.1063/0.0506	35.248/37.408	0.972105/0.982213	
		0.0203/0.0060	39.799/42.366	0.989180/0.993567	
Lung	2D method	0.0046/0.0029	45.222/49.363	0.996444/0.998398	0.049/24.841
		0.0168/0.0026	36.198/38.479	0.977824/0.986030	
		0.0065/0.0046	40.933/43.528	0.991438/0.994828	
Lung	Proposed	0.0006/0.0008	46.290/49.896	0.997044/0.998568	0.105/8.522

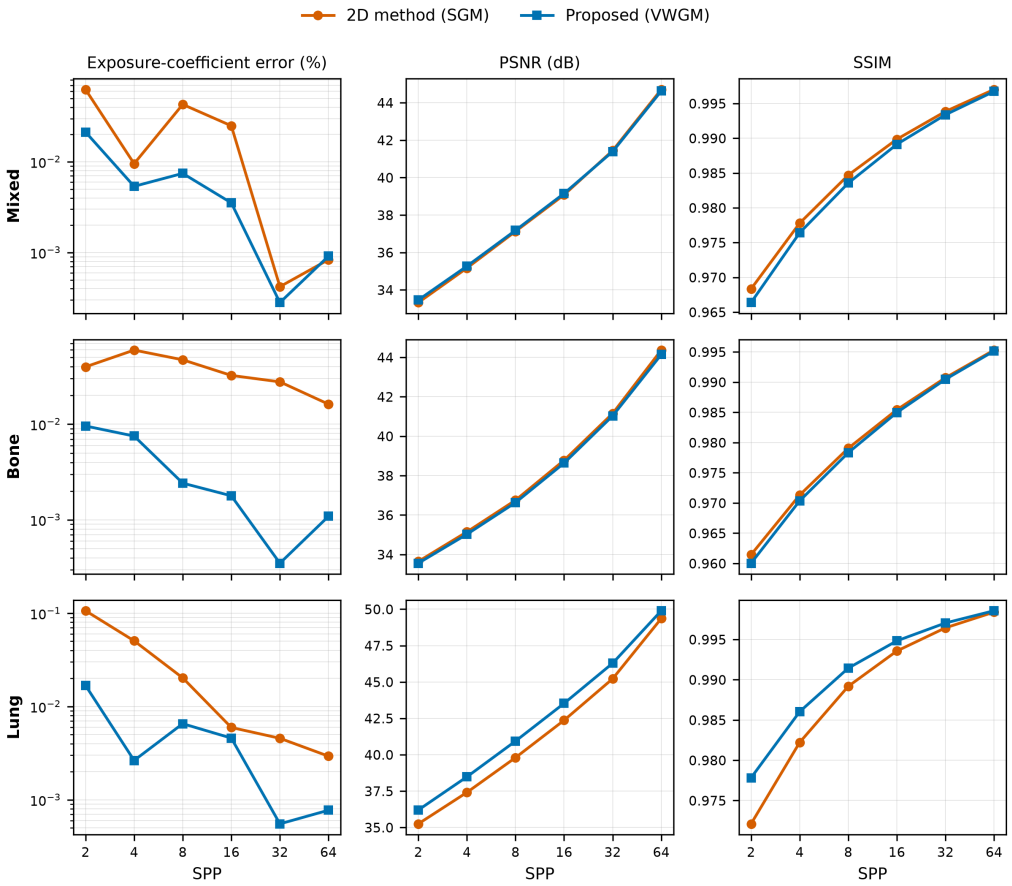


Figure 3. SPP-dependent convergence of the two-dimensional screen-space method and the proposed visibility-weighted volumetric method. Errors are computed relative to each method's own 128-SPP result.

Across the six non-reference SPP levels, the mean relative coefficient error is 0.005% for VWGM and 0.031% for SGM; the maxima are 0.021% and 0.106%. PSNR and SSIM are similar for mixed and bone rendering, whereas VWGM is higher at every tested SPP for lung rendering. The volumetric estimate is therefore less sensitive to sampling variation.

4.3. Target-Gray Sensitivity

We vary G_a over $\{0.18, 0.35, 0.50, 0.65\}$ for the same CT volume under the mixed, bone, and lung transfer functions. All results use 128 SPP with the camera, illumination, and remaining rendering settings fixed. For each transfer function, the same 160×160-pixel ROI is retained at every target gray. ROI contrast is the normalized RMS contrast of the non-black ROI luminance values, $C_{\mathrm{ROI}} = \sigma_Y / \mu_Y$. Table 5 and Figure 4 report the resulting metrics and the corresponding images.

Table 5. Target-gray sensitivity at 128 SPP. Values are the mean ± sample standard deviation across the mixed, bone, and lung transfer functions.

Target	δ	Mean Luminance	C_{ROI}	Sat. (%)
$G_a = 0.18$	0.633 ± 0.133	48.494 ± 8.103	0.471 ± 0.143	0.077 ± 0.061
$G_a = 0.35$	1.231 ± 0.259	93.242 ± 15.041	0.458 ± 0.136	3.040 ± 2.168
$G_a = 0.50$	1.759 ± 0.370	127.601 ± 18.623	0.412 ± 0.139	15.015 ± 6.387
$G_a = 0.65$	2.287 ± 0.480	153.332 ± 19.654	0.343 ± 0.144	30.661 ± 7.259

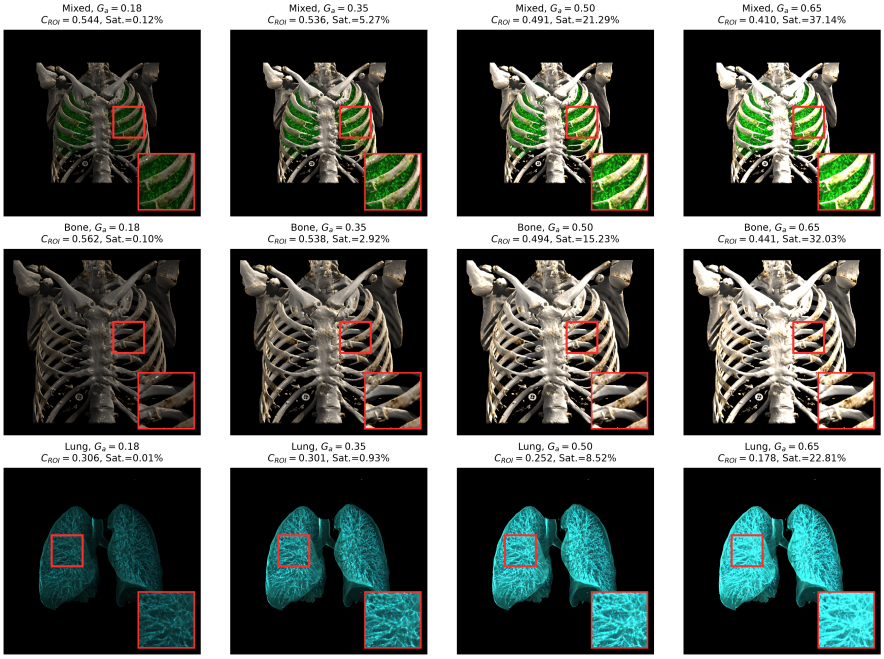


Figure 4. Target-gray sensitivity under the mixed, bone, and lung transfer functions. Red boxes identify the fixed ROI for each transfer function, and the lower-right insets show the corresponding magnified regions.

Increasing G_a raises the exposure coefficient and output luminance but also increases highlight saturation and reduces ROI contrast. We therefore use $G_a = 0.50$ as a practical compromise for the reported experiments; it is not presented as a universal optimum.

4.4. Numerical-Stabilization Sensitivity

We test $\epsilon \in \{10^{-8}, 10^{-7}, 10^{-6}, 10^{-5}, 10^{-4}\}$ at 128 SPP and $G_a = 0.50$. Alongside the complete volumes, the mixed case uses a partial crop, the bone case a zero-valued exterior, and the lung case both perturbations. Each configuration is compared with its own $\epsilon = 10^{-6}$ result. For any quantity x, the relative change is $E_x(\epsilon) = |x_\epsilon - x_{10^{-6}}|/|x_{10^{-6}}| \times 100\%$. Table 6 and Figure 5 summarize the worst-case deviation of each configuration.

Table 6. Numerical-stability summary for the complete and low-support configurations. Each entry is the worst case over the non-reference ϵ values relative to that configuration's $\epsilon = 10^{-6}$ result.

TF	Configuration	Max. E_{μ_g} (%)	Max. E_δ (%)	Max. $E_{\bar{Y}}$ (%)	Min. SSIM	Max. $\lvert\Delta$Sat.$\rvert$ (pp)
Mixed	Full	0.0234	0.0234	0.0150	0.999943	0.0208
Mixed	Partial crop	0.0299	0.0299	0.0195	0.999972	0.0093
Bone	Full	0.0175	0.0175	0.0140	0.999904	0.0050
Bone	Zero-valued exterior	0.0175	0.0175	0.0478	0.999984	0.0164
Lung	Full	0.0204	0.0204	0.0181	0.999903	0.0107
Lung	Crop + zero exterior	0.0183	0.0183	0.0189	0.999954	0.0169

All six configurations are insensitive to the tested range of ϵ. The largest changes in exposure coefficient and final mean luminance are below 0.03% and 0.05%, respectively; the minimum SSIM exceeds 0.9999, and the saturation ratio changes by less than 0.021 percentage points. The default $\epsilon = 10^{-6}$ is therefore retained for the remaining experiments.

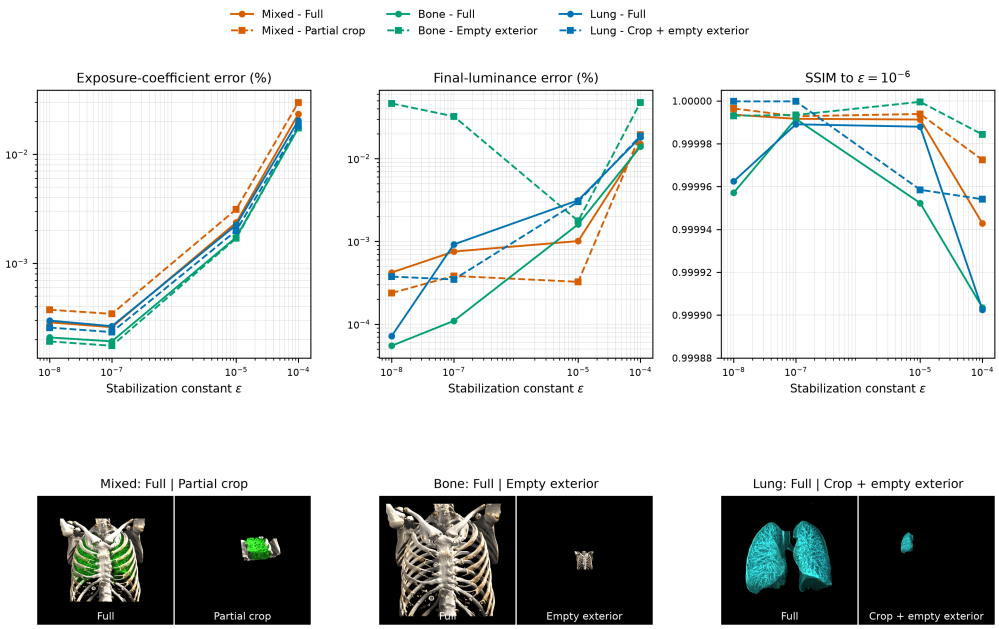


Figure 5. Numerical-stabilization sensitivity for the complete and low-support configurations. The upper row shows exposure-coefficient error, final-luminance error, and SSIM relative to each configuration's default $\epsilon = 10^{-6}$ result. The lower row shows the complete and perturbed configurations at the default ϵ.

4.5. Averaging-Operator Ablation

The averaging-operator ablation keeps the transport path, effective samples, and weights $w_i = T_i \alpha_i$ unchanged. VWAM uses a weighted arithmetic mean, whereas VWGM uses a weighted geometric mean; both produce one global coefficient.

For a direct component decomposition, let L_p denote the luminance of the P nonzero composited pixels and let $\bar{w}_i = w_i / \sum_k w_k$. The four estimators can be written as

$$\mu_{\mathrm{SGM}} + \epsilon = \exp\left(\frac{1}{P} \sum_{p=1}^{P} \log(L_p + \epsilon) \right),$$

$$\mu_{\mathrm{VUGM}} + \epsilon = \exp\left(\frac{1}{K} \sum_{i=1}^{K} \log(\mathcal{E}_i + \epsilon) \right),$$

$$\mu_{\mathrm{VWGM}} + \epsilon = \exp\left(\sum_{i=1}^{K} \bar{w}_i \log(\mathcal{E}_i + \epsilon) \right),$$

$$\mu_{\mathrm{VWAM}} = \sum_{i=1}^{K} \bar{w}_i \mathcal{E}_i.$$

The paired substitutions SGM→VUGM, VUGM→VWGM, and VWAM→VWGM change only the statistical domain, visibility weights, and averaging operator, respectively, thereby isolating the three effects.

Each operator is compared with its own 128-SPP result. We report the mean coefficient deviation over 2–64 SPP, the output-luminance CV over 2–128 SPP, SSIM at 64 SPP, and the mean luminance and saturation ratio at 128 SPP. All other parameters are fixed. Figure 6 compares the four operators visually.

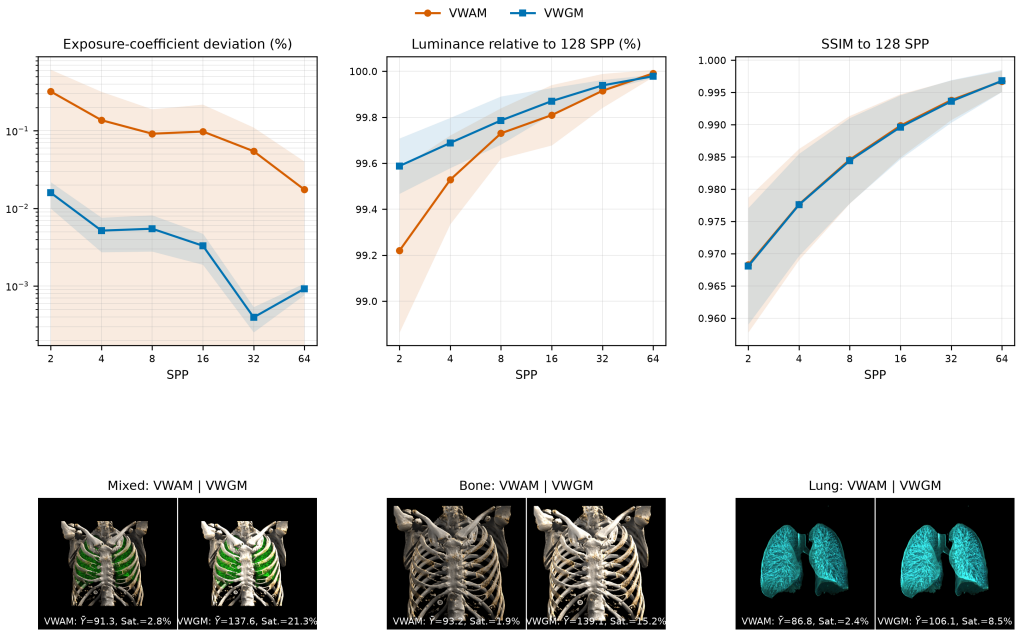


Figure 6. Averaging-operator ablation. The upper row shows the mean over the three transfer functions for exposure-coefficient deviation, output mean luminance, and SSIM relative to each operator's own 128-SPP reference. The orange and blue curves denote VWAM and VWGM, respectively; the corresponding translucent bands indicate ± one sample standard deviation across the three transfer functions. The lower row compares the 128-SPP VWAM and VWGM results; annotations give output mean luminance and saturation.

VWGM reduces the mean coefficient deviation from 0.120% to 0.005% and the luminance CV from 0.264% to 0.145%, while the two operators have similar mean SSIM at 64 SPP. Its stronger response to low-energy samples also raises the mean saturation ratio at 128 SPP from 2.37% to 15.02%. Thus, the gain in sampling stability comes with a higher risk of highlight clipping.

To separate these effects, we additionally evaluated an unweighted volumetric geometric mean (VUGM). SGM and VUGM use the same geometric-mean form; thus, the reduction in mean exposure-coefficient deviation from 0.0309% to 0.0020% quantifies the overall effect of moving the statistic from composited screen pixels to pre-projection volumetric samples. VUGM and VWGM use the same positive volumetric samples and geometric-mean form; adding the $T_i\alpha_i$ visibility weights reduces the mean 128-SPP saturation ratio across the three transfer functions from 36.75% to 15.02%, a reduction of 21.73 percentage points, by limiting the disproportionate influence of weak, low-visibility samples on the global exposure coefficient. Finally, VWAM and VWGM in Table 7 use the same volumetric samples and weights, so the changes in coefficient deviation from 0.120% to 0.005% and luminance CV from 0.264% to 0.145% are attributable to replacing the arithmetic mean with the geometric mean. Accordingly, volumetric statistics account for most of the coefficient-deviation reduction, visibility weighting primarily limits overexposure caused by low-visibility samples, and the averaging form determines the remaining trade-off between stability and saturation. Figure 7 shows the corresponding VUGM and VWGM images at 128 SPP.

Table 7. Averaging-operator ablation. Lower is better for exposure-coefficient deviation and luminance CV. Rows shown in bold are summary rows: they report the mean ± sample standard deviation over the three transfer functions, and the bold type distinguishes them from the individual per-transfer-function rows listed above.

Variant	TF	Coeff. Dev. (%)	Lum. CV (%)	$SSIM_{64}$	$\bar{Y}_{128}$	$Sat._{128}$ (%)
VWAM	Mixed	0.2476	0.3389	0.996128	91.29	2.812
VWAM	Bone	0.1070	0.3249	0.995530	93.16	1.895
VWAM	Lung	0.0051	0.1284	0.998602	86.82	2.412
VWAM mean	—	0.120 ± 0.122	0.264 ± 0.118	0.9968 ± 0.0016	90.4 ± 3.3	2.37 ± 0.46
VWGM	Mixed	0.0065	0.1345	0.996734	137.60	21.290
VWGM	Bone	0.0038	0.1944	0.995145	139.09	15.233
VWGM	Lung	0.0053	0.1045	0.998568	106.11	8.522
VWGM mean	—	0.0052 ± 0.0013	0.145 ± 0.046	0.9968 ± 0.0017	127.6 ± 18.6	15.02 ± 6.39

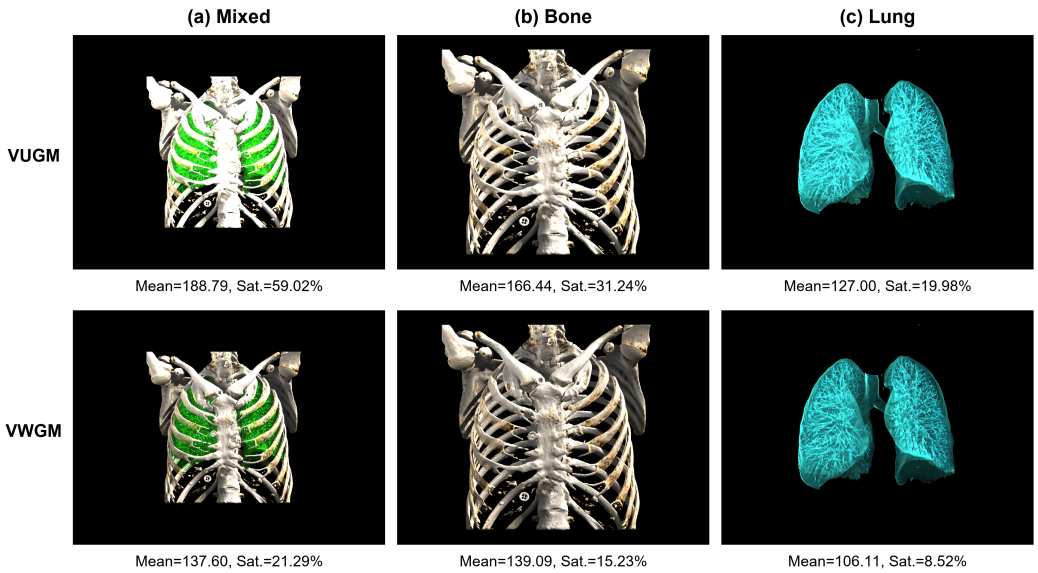


Figure 7. Visibility-weighting ablation at 128 SPP. The top row shows the unweighted volumetric geometric mean (VUGM), and the bottom row shows the visibility-weighted geometric mean (VWGM); the columns show the mixed, bone, and lung transfer functions.

4.6. Illumination and Transfer-Function Robustness

We vary the photon energy of all three lights by the common scale $s_E \in \{0.8, 1.0, 1.3\}$ and keep the camera, transfer function, 128-SPP sampling, and other settings fixed. Under uniform energy scaling, μ_g should vary with s_E and δ inversely with s_E; hence, $s_E\delta$ should remain nearly constant. This relation is first checked for the mixed transfer function. Table 8 reports this verification for the mixed transfer function.

Table 8. Verification of the predicted linear photon-energy scaling law for VWGM under the mixed transfer function.

s_E	μ_g	δ	$s_E\delta$	Final Mean	$SSIM_{1.0}$	Saturation (%)
0.8	0.238270	2.098461	1.678769	136.040	0.998541	20.478
1.0	0.294369	1.698547	1.698547	137.598	1.000000	21.290
1.3	0.377465	1.324624	1.722012	139.374	0.998224	22.212

The same energy sweep is then applied to all three transfer functions. Table 9 reports luminance and ROI-contrast CV, minimum SSIM, maximum normalized gradient RMSE relative to the default-energy frame, and maximum saturation.

Table 9. Robustness of VWGM to illumination changes under the three transfer functions and relative energy scales $s_E \in \{0.8, 1.0, 1.3\}$.

Transfer Function	Lum. CV (%)	Min. SSIM	ROI CV (%)	Max. Grad. NRMSE	Max. Sat. (%)
Mixed	0.989	0.998224	0.961	0.0306	22.212
Bone	1.090	0.997931	0.495	0.0330	16.179
Lung	1.098	0.998610	1.660	0.0395	9.075

Across the three transfer functions, $s_E\delta$ differs from its default-energy value by at most 1.38%. The final-luminance CV is 0.99–1.10%, the minimum SSIM is 0.998, and the ROI-contrast CV is at most 1.66%. The exposure and local structure are therefore stable over the tested illumination range. Figure 8 shows the corresponding rendered images and the fixed ROIs.

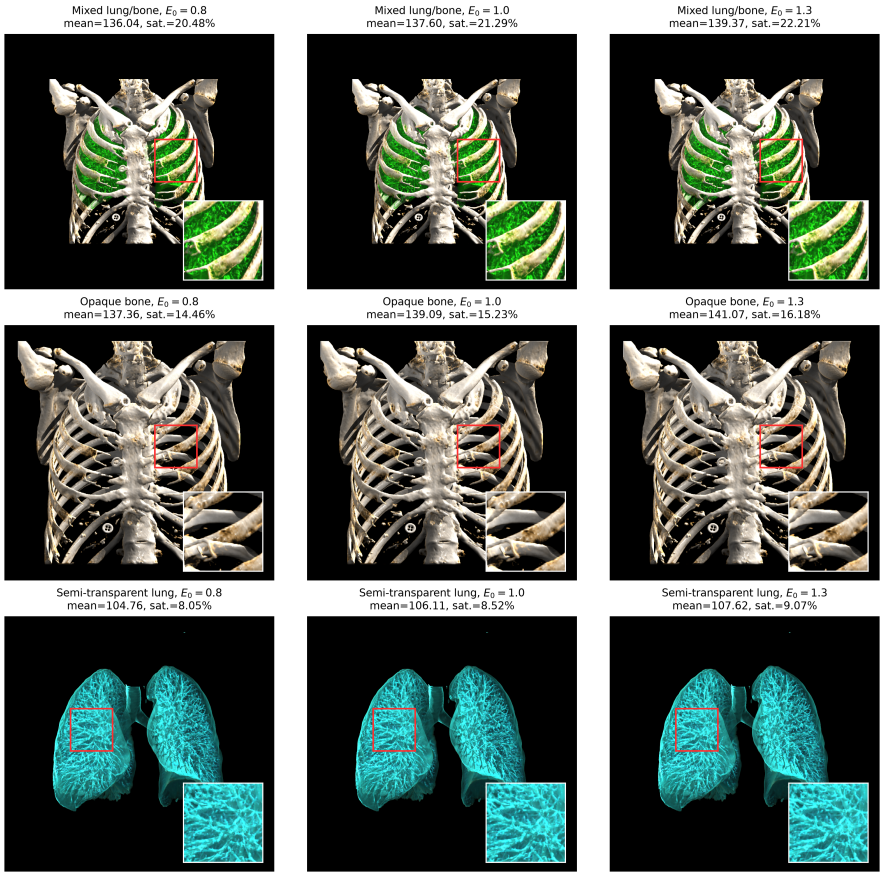


Figure 8. Proposed-method outputs under the three transfer functions and three relative photon-energy scales. Red boxes identify the fixed ROI reused for each transfer function, and insets show the corresponding magnified regions.

4.7. Comparison with Two-Dimensional Methods and Structural Detail

Durand–Dorsey two-dimensional HDR tone mapping is compared with the proposed method for the mixed, bone, and lung transfer functions of the lung CT dataset.

Its exposure-compensation parameter is varied to produce dark, balanced, and over-exposed images. Structural detail is measured in matched ROIs. Table 10 reports these metrics, and Figure 9 shows the compared images together with the matched ROIs.

Table 10. Structural-detail metrics within the magnified ROIs for the high-exposure two-dimensional HDR result and the proposed method at 128 SPP, reported as mean ± standard deviation over the three transfer functions.

Method	ROI RMS Contrast	ROI Normalized Gradient
2D HDR (0 EV)	0.254 ± 0.051	0.087 ± 0.031
Proposed	0.628 ± 0.268	0.160 ± 0.061

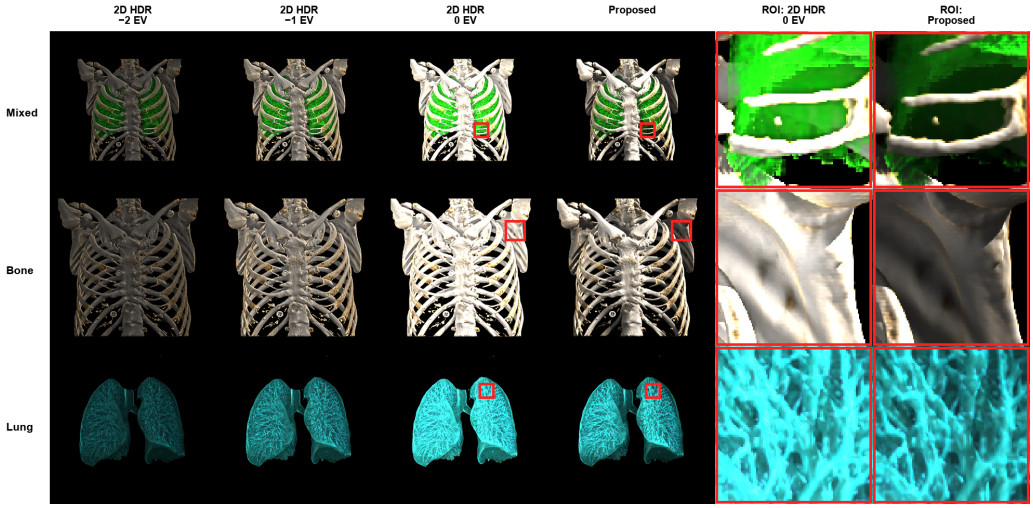


Figure 9. The three exposure settings of two-dimensional HDR and the 128-SPP proposed-method output under the three transfer functions. Red boxes mark the ROI fixed within each transfer function and reused across methods; the two rightmost columns magnify the same locations for the high-exposure two-dimensional result and the proposed result. The mixed and lung ROIs emphasize depth structure within semi-transparent tissue.

The mean normalized ROI gradient is 0.087 for the high-exposure two-dimensional result and 0.160 for the proposed method. In the mixed and lung views, the matched enlargements retain clearer vascular branches, tissue boundaries, and depth-dependent intensity changes with the proposed method.

4.8. Computational Cost and Memory

Let P be the number of output pixels, S the SPP, M the average samples visited per ray, and L_s the lights evaluated per sample. Baseline transport costs $O(PSML_s)$. Exposure estimation adds $O(PSM)$ work for logarithms, weighted accumulation, and reduction, followed by $O(P)$ output scaling. The total cost is $O(PSML_s + PSM + P)$.

Each retained sample stores a three-dimensional coordinate and a depth value in one `float4` (16 bytes). For N_r resident samples, the buffer size is therefore $16N_r$ bytes. The tested implementation uses 2 SPP and records about 20–200 effective samples per pixel path, depending on the transfer function. Storage for both passes is preallocated at the 200-sample upper bound. Table 11 summarizes the resulting auxiliary storage requirement.

Table 11. Theoretical auxiliary GPU-memory requirement of the proposed implementation at a 512×512 output resolution, 2 SPP, and 20–200 effective samples per pixel path.

Item	Storage Expression	Memory at 512^2
Coordinate–depth record	`float3` + `float` = 16 bytes	16 bytes/sample
Valid sample records at 2 SPP	$512^2 \times 2 \times [20, 200] \times 16$ bytes	160–1600 MiB
Preallocated sample buffer	$512^2 \times 2 \times 200 \times 16$ bytes	1600 MiB (1.563 GiB)
Total preallocated auxiliary storage	$16N_r$ bytes	1600 MiB (1.563 GiB)

At 512 × 512 and 2 SPP, the observed range of 20–200 samples corresponds to 160–1600 MiB of records. The buffer is preallocated at the upper bound and therefore occupies 1600 MiB (1.563 GiB); transfer functions change the number of records written, not this capacity. The observed sample counts and memory use agree with the estimate. Because stored depth records grow linearly with SPP, high-SPP runs process them in segments and merge the partial statistics to bound peak memory.

4.9. Cross-Dataset Evaluation

The cross-dataset experiment renders three volumes from the open-source CPTAC dataset [17], together with a cardiac CT volume, an MR volume, and a PT volume. Figure 10 shows the rendering results, and Table 12 lists the corresponding volume dimensions and rendering time.

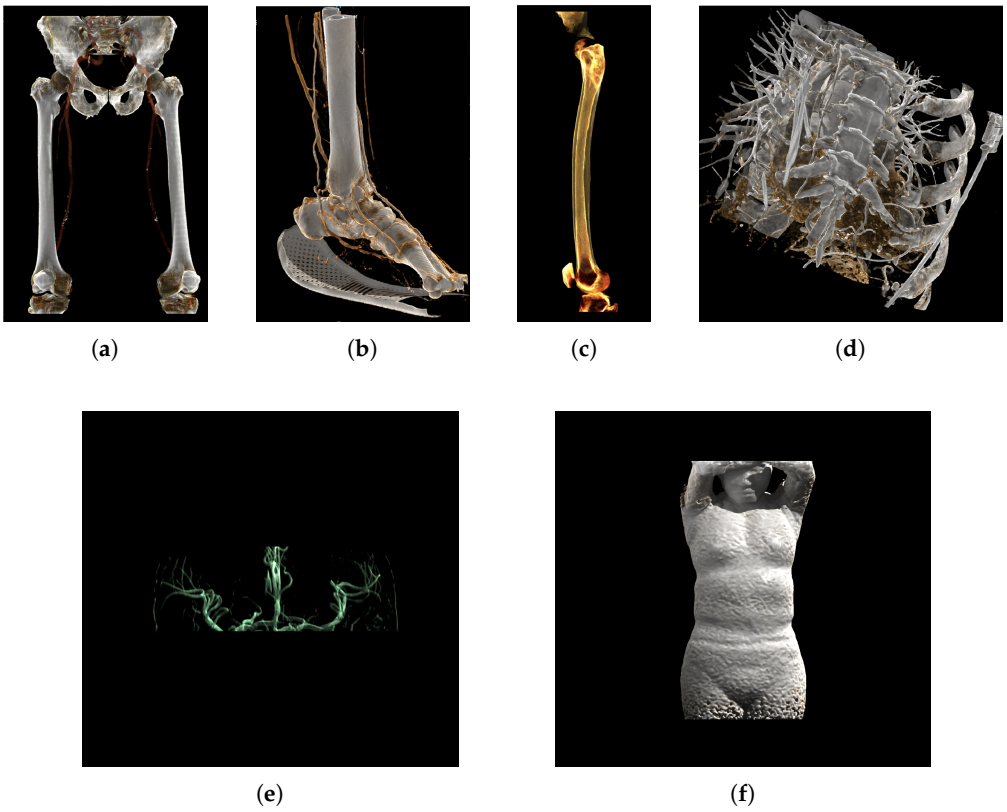


Figure 10. Rendering results on six medical volume datasets. (**a**) C3N-00843. (**b**) C3L-01466. (**c**) C3L-03177. (**d**) Cardiac. (**e**) MR. (**f**) PT.

Table 12. Volume dimensions and rendering time in the cross-dataset evaluation.

Dataset	Volume Size	Rendering Time (ms)
C3N-00843	512 × 512 × 736	75
C3L-01466	512 × 512 × 216	68
C3L-03177	512 × 512 × 545	62
Cardiac	512 × 512 × 305	84
MR	512 × 512 × 67	56
PT	256 × 256 × 405	81

The transfer function maps each dataset's scalar values to color and opacity: CT commonly uses Hounsfield units, whereas MR and PT use their own intensity scales. Exposure is estimated from the resulting sample energy, opacity, and visibility rather than from a modality label or a fixed raw-intensity unit. The same procedure can therefore be

used after a suitable transfer function has been defined, although the final appearance still depends on both the data distribution and that mapping.

5. Discussion

The specific contribution is not a new local tone-mapping operator. Instead, the global exposure scale is estimated from visibility-weighted effective volumetric-sample energies before projection rather than from projected screen pixels. The estimated scale is applied uniformly to all accumulated colors and does not introduce two-dimensional local compression across image regions.

Like Reinhard-style global tone reproduction, the proposed method uses a logarithmic-domain statistic to estimate one uniform global scale. The difference is that Reinhard-style estimation is performed on projected two-dimensional luminance, whereas our statistic is computed before projection and weighted by transmittance and opacity [13]. Unlike the local Durand–Dorsey operator, our method neither decomposes the two-dimensional image into base and detail layers nor applies spatially varying compression [14]. Such a local operator is, in principle, better suited to spatial exposure conflicts, whereas our method retains uniform scaling and the associated volumetric-contribution relationships.

Dinesha et al. use HDR volume rendering to represent uncertainty in volumetric data, with emphasis on displaying a large dynamic range [15]. Xu et al. study how to maintain volume-rendering quality under dynamic illumination, with emphasis on the rendered appearance as lighting changes [16]. Neither computes a global exposure coefficient from the transmittance, opacity, and sample energy accumulated during volumetric integration. Our method directly uses this information, already produced during integration, to estimate global exposure before projection.

The method is intended for the conventional medical-volume-rendering configuration evaluated here, in which multiple lights provide broadly balanced overall illumination so that major structures can be viewed together. Within this operating range, one global coefficient can be used directly for automatic exposure adjustment in volume rendering. Its principal limitation is that it does not handle extremely intense illumination, localized extreme highlights, or very large spatial brightness differences. In such cases, one coefficient cannot accommodate bright and dark regions simultaneously, and clipping or exposure conflicts may occur; region-aware exposure or local tone mapping would then be required.

6. Conclusions

We presented a voxel-contribution-based adaptive exposure model for physically based volume rendering, using a visibility-weighted geometric mean of volumetric-sample energies to estimate one global exposure coefficient. Compared with the screen-space geometric mean, the mean exposure-coefficient deviation decreases from 0.031% to 0.005%; across the tested global-illumination range, output-intensity variation remains below 1.10% and the minimum SSIM is 0.998. The ablation further shows that the volumetric domain, visibility weighting, and averaging form affect exposure stability, overexposure control, and the trade-off between stability and saturation, respectively. The same procedure is applied to CT, MR, and PT volumes and is intended for conventional medical volume rendering with broadly balanced multi-light illumination, where it can be used directly for automatic exposure adjustment. Because it remains a global adjustment, it cannot resolve exposure conflicts caused by extremely intense illumination, localized extreme highlights, or very large spatial brightness differences.

7. Future Work

Future work will extend established two-dimensional HDR and tone-mapping strategies for extreme illumination and local exposure conflicts to the three-dimensional volumetric domain. In practice, histogram analysis and local exposure control will be moved from the two-dimensional image to the pre-projection volumetric domain, where volumetric energy and visibility will determine the target gray and regional exposure coefficients for extreme or localized illumination.

Author Contributions: Conceptualization, Y.Y., J.Y., Q.S. and Y.H.; Validation, J.Y.; Formal analysis, Y.Y.; Investigation, Y.Y., Q.S. and Y.H.; Data curation, Y.Y. and J.Y.; Writing—original draft, Q.S.; Writing—review and editing, Y.Y. and J.Y.; Visualization, Y.H.; Supervision, J.Y.; Project administration, J.Y.; Funding acquisition, Y.Y. All authors have read and agreed to the published version of the manuscript.

Funding: The authors received no specific funding for this work.

Institutional Review Board Statement: Ethical review and approval were waived. This computer graphics study used only pre-existing, fully de-identified volumetric scanning data for the technical evaluation of a rendering algorithm. No participant recruitment, clinical intervention, new medical data acquisition, identifiable personal information, or patient-level clinical analysis was involved.

Informed Consent Statement: Informed Consent was waived. This computer graphics study used only pre-existing, fully de-identified volumetric scanning data for the technical evaluation of a rendering algorithm. No participant recruitment, clinical intervention, new medical data acquisition, identifiable personal information, or patient-level clinical analysis was involved.

Data Availability Statement: The data presented in this study are openly available in [CPTAC-SAR] at [http://doi.org/10.7937/tcia.2019.9bt23r95].

Conflicts of Interest: The authors declare no competing interests.

References

1. Salama, C.R. Gpu-based monte-carlo volume raycasting. In *Proceedings of the 15th Pacific Conference on Computer Graphics and Applications (PG'07)*; IEEE: New York, NY, USA, 2007; pp. 411–414. https://doi.org/10.1109/PG.2007.27.
2. Westover, L.A. *Splatting: A Parallel, Feed-Forward Volume Rendering Algorithm*; The University of North Carolina at Chapel Hill: Chapel Hill, NC, USA, 1991.
3. Lacroute, P.; Levoy, M. Fast volume rendering using a shear-warp factorization of the viewing transformation. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*; ACM: New York, NY, USA, 1994; pp. 451–458. https://doi.org/10.1145/192161.192283.
4. Bishop, G.; Weimer, D.M. Fast phong shading. *ACM Siggraph Comput. Graph.* **1986**, *20*, 103–106. https://doi.org/10.1145/15886.15897.
5. Lafortune, E.P.; Willems, Y.D. Bi-directional path tracing. In Proceedings of the Third International Conference on Computational Graphics and Visualization Techniques (Compugraphics '93), Alvor, Portugal, 6–10 December 1993; pp. 145–153. Available online: https://graphics.cs.kuleuven.be/publications/BDPT/ (accessed on 25 October 2026).
6. Veach, E.; Guibas, L. Bidirectional estimators for light transport. In *Photorealistic Rendering Techniques*; Springer: Cham, Switzerland, 1995; pp. 145–167. https://doi.org/10.1007/978-3-642-87825-1_11.
7. Yuan, L.; Sun, J. Automatic exposure correction of consumer photographs. In Proceedings of the Computer Vision–ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, 7–13 October 2012; Proceedings, Part IV 12; Springer: Cham, Switzerland, 2012; pp. 771–785. https://doi.org/10.1007/978-3-642-33765-9_55.
8. Reinhard, E. High dynamic range imaging. In *Computer Vision: A Reference Guide*; Springer: Cham, Switzerland, 2020; pp. 1–6. https://doi.org/10.1007/978-3-030-03243-2_843-1.
9. Csebfalvi, B.; Szirmay-Kalos, L. Monte Carlo volume rendering. In *Proceedings of the IEEE Visualization 2003, VIS 2003*; IEEE: New York, NY, USA, 2003; pp. 449–456. https://doi.org/10.1109/VISUAL.2003.1250406.
10. Csebfalvi, B. Interactive transfer function control for Monte Carlo volume rendering. In *Proceedings of the 2004 IEEE Symposium on Volume Visualization and Graphics*; IEEE: New York, NY, USA, 2004; pp. 33–38. https://doi.org/10.1109/SVVG.2004.6.

11. Zhang, Y.; Dong, Z.; Ma, K.L. Real-time volume rendering in dynamic lighting environments using precomputed photon mapping. *IEEE Trans. Vis. Comput. Graph.* **2013**, *19*, 1317–1330. https://doi.org/10.1109/TVCG.2013.17.
12. Iglesias-Guitian, J.A.; Mane, P.S.; Moon, B. Real-time denoising of volumetric path tracing for direct volume rendering. *IEEE Trans. Vis. Comput. Graph.* **2020**, *28*, 2734–2747. https://doi.org/10.1109/TVCG.2020.3037680.
13. Reinhard, E.; Stark, M.; Shirley, P.; Ferwerda, J. Photographic tone reproduction for digital images. In *Seminal Graphics Papers: Pushing the Boundaries*; ACM: New York, NY, USA, 2023; Volume 2, pp. 661–670. https://doi.org/10.1145/566654.566575.
14. Durand, F.; Dorsey, J. Fast bilateral filtering for the display of high-dynamic-range images. *ACM Trans. Graph.* **2002**, *21*, 257–266. https://doi.org/10.1145/566654.566574.
15. Dinesha, V.; Adabala, N.; Natarajan, V. Uncertainty visualization using HDR volume rendering. *Vis. Comput.* **2012**, *28*, 265–278. https://doi.org/10.1007/s00371-011-0614-7.
16. 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 Trans. Vis. Comput. Graph.* **2024**, *31*, 5288–5303. https://doi.org/10.1109/TVCG.2024.3445339.
17. National Cancer Institute Clinical Proteomic Tumor Analysis Consortium (CPTAC). *The Clinical Proteomic Tumor Analysis Consortium Sarcomas Collection (CPTAC-SAR) (Version 10) [Data set]*; The Cancer Imaging Archive: Frederick, MD, USA, 2019. https://doi.org/10.7937/tcia.2019.9bt23r95.

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.