

# Real-time Volume Rendering in Dynamic Lighting Environments using Precomputed Photon Mapping

Yubo Zhang, *Student Member, IEEE*, Zhao Dong, *Member, IEEE*, and Kwan-Liu Ma, *Fellow, IEEE*

**Abstract**—We present a framework for precomputed volume radiance transfer that achieves real-time rendering of global illumination effects for volume datasets such as multiple scattering, volumetric shadows, etc. Our approach incorporates the volumetric photon mapping method into the classical precomputed radiance transfer pipeline. We contribute several techniques for light approximation, radiance transfer precomputation and real-time radiance estimation, which are essential to make the approach practical and to achieve high frame rates. For light approximation, we propose a new discrete spherical function which has better performance for construction and evaluation when compared with existing rotational invariant spherical functions such as spherical harmonics and spherical radial basis functions. In addition, we present a fast splatting-based radiance transfer precomputation method and an early-evaluation technique for real-time radiance estimation in the clustered PCA space. Our techniques are validated through comprehensive evaluations and rendering tests. We also apply our rendering approach to volume visualization.

**Index Terms**—Volume Rendering, Precomputed Radiance Transfer, Volume Ray Casting, Multiple Scattering, Volume Shadow.

## 1 INTRODUCTION

Light traversing a volume usually undergoes absorption and scattering events. Before reaching the view point, the radiance energy can scatter one time (i.e. single scattering (SS)) or multiple times (i.e. multiple scattering (MS)) inside the volume. Following the definition of surface light transport, we can regard the shading from SS as the *direct illumination* and from MS as the *indirect illumination* or *global illumination* (GI) for the volume dataset. The volume GI effects can greatly enhance the visual perception of the volume dataset. For example, global occlusion and shadow effects present better depth order information and MS effects can also highlight the local thickness of the spatial structures given appropriate light settings. However, simulating all light transportation within volume including absorption and scattering is time consuming. Therefore, most existing methods for volume rendering have to strike a balance between performance and quality. Great efforts have been made to accelerate the light simulation process while preserving shading quality such as two-step radiance estimation [1] and volumetric photon mapping [2]. Although they achieve high quality results, it is difficult to achieve interactive frame rates with these techniques. For comparison, classic volume ray casting with simple Phong shading can achieve real-time performance using modern graphics processing units (GPUs). Still, even with

advanced shading techniques such as ambient occlusion and approximated scattering effects, the quality and accuracy of the real-time methods cannot match the previously mentioned offline rendering techniques. To achieve real-time performance and high visual quality at the same time, a practical strategy is to introduce preprocessing.

Recently, the precomputed radiance transfer (PRT) technique was developed for real-time surface rendering with GI effects including soft shadows and subsurface scattering [3]. Its application for rendering surface meshes gives promising results, so many extensions and improvements have been made. However, most PRT-based research focuses on rendering surface meshes, and the design of a general PRT framework for volume rendering has not received much attention. The issues involved are lighting approximation, volume radiance approximation, efficient precomputation of general volumetric effects, volume radiance compression, and sampling strategy during real-time volume rendering. To achieve efficient precomputation while maintaining volume GI effects, volumetric photon mapping (VPM) [2] is a practical method to use. The major challenges are the seamless incorporation of VPM into the precomputation stage with an appropriate photon casting method and the determination of a fast radiance estimation strategy. We present a novel precomputed volume radiance transfer (PVRT) method to achieve successful integration so that high-quality volume GI effects can be rendered in real-time (Fig.1). With PVRT, we introduce a photon casting method that generates all the photons at the volume boundary. Each photon has varying energy and may also carry negative energy. To estimate the radiance, we adopt an efficient splatting-based radiance estimation method. We also contribute a new discrete spherical function (DSF) for lighting approximation. DSF achieves better quality than

- Yubo Zhang and Kwan-Liu Ma are with the Department of Computer Science, University of California Davis, One Shields Avenue, Davis, CA 95616.  
E-mail: {ybzhang,klma}@ucdavis.edu
- Zhao Dong is with the Program of Computer Graphics at Cornell University, 588 Rhodes Hall, Cornell University, Ithaca, NY 14853.  
E-mail: zd@graphics.cornell.edu

the classic spherical harmonics (SH) basis function for approximating high-frequency lighting, and unlike wavelet basis functions, it is also rotationally invariant. In addition, the interpolation property of DSF enables much faster construction and evaluation when compared with spherical radial basis functions (SRBF) [4], while the approximation quality is similar. To summarize, our contributions include:

- A novel DSF for light approximation, which is rotationally invariant, fast to construct and evaluate, and has the ability to approximate high-frequency lighting.
- A PVRT method that successfully incorporates the VPM method into the PRT framework and achieves efficient and general volume radiance estimation.
- A comprehensive performance and quality evaluation of our approach. We also demonstrate its application in interactive scientific and medical volume visualization.

The remainder of this paper is organized as follows. Sec.2 describes existing works and Sec.3 reviews the background of light transport in participating media. After an overview of the PVRT method in Sec.4. Sec.5 introduces the new discrete spherical function (DSF) and its properties. Next, Sec.6 and Sec.7 discuss the general procedure for incorporating the VPM technique into the PRT framework as well as the final rendering step. Several implementation details are discussed in Sec.8. The experimental results and a comprehensive evaluation of our method are given in Sec.9. Finally, Sec.10 concludes the paper and introduces plan for future works.

## 2 RELATED WORKS

We refer readers to [5], [6] for a complete survey of volume GI rendering, and to [7], [8] for a comprehensive introduction of precomputed radiance transfer (PRT) methods. In this section, we only cover the most relevant works.

**Offline Volume GI Rendering** The rendering of volume GI effects is essentially the solution of the radiative transfer equation [9]. Kajiya and Von Herzen [1] presented a two-pass ray tracing method. First, to estimate radiance for each voxel, the radiative transfer equation is approximately evaluated using spherical harmonics (SH). Thereafter, ray marching is performed along view rays to gather the final radiance. Since the MS inside the volume usually exhibits low-frequency behavior, only the first few SH bases are necessary for generating visually convincing results. Afterwards, several methods have been proposed to compute numerical solutions of the radiative transfer equation by using either ray tracing [10], [11] or radiosity [12]. These solutions can render various volume GI effects, but the computational cost is enormous, so they can only run in offline mode. To improve performance, Stam [13] suggested to approximate MS effects using the diffusion process, which can be simulated by either a multi-grid scheme or a finite-element blob method. Nevertheless, the speed is still far from interactive.

The volumetric photon mapping (VPM) method [2] was proposed to render volume caustics. In its first stage, photons are traced through the volume and stored into a hierarchical data structure, like a kd-tree. Afterward,

ray marching is performed and at each step the stored photons are gathered to compute the final radiance. This final radiance gathering step can also be computed more efficiently by photon splatting [14]. The creation and maintenance of the hierarchical data structure prevents VPM from achieving interactive performance. Even with a GPU-based kd-tree [15], rendering high-frequency volume caustics with tens of millions of photons is still non-trivial. Our PVRT method also adopts photon tracing to simulate light transport, and efficiently precomputes the radiant exitance for each voxel based on photon splatting.

**Basis Function in PRT** In the precomputation stage, PRT-based methods simulate complex light transfer and approximate the radiance transfer using certain kinds of basis functions. Then, the radiant exitance can be approximately reconstructed by multiplying the coefficients of incident lighting and the radiance transfer matrix. Several functions have been successfully applied in PRT, such as SH [3], wavelets [16], [17], spherical radial basis functions (SRBF) [4], Gaussians [18], etc. Each of these functions has its own pros and cons. The SH basis smoothly reconstructs low-frequency lighting effects very well and can be freely rotated, but it cannot efficiently represent high-frequency signals. The wavelet basis can represent all-frequency effects with a small number of coefficients, but its rotation is non-trivial and could introduce temporal flickering due to nonlinear reconstruction. The SRBF can reconstruct all-frequency effects and is easy to rotate, but its precomputation cost is much higher than that of other functions. To realize all-frequency rendering, we introduce a novel discrete spherical function (DSF) that achieves similar quality to SRBF with much less precomputation.

**Precomputation-based Volume GI Rendering** Most existing PRT-based methods deal with the light transfer between rigid surfaces, and recently some advanced techniques can additionally handle deformable scenes [19]. However, precomputation-based GI rendering for volume datasets has received far less attention. Wyman et al. [20] introduced a method to precompute the GI using SH for interactive rendering of isosurfaces extracted from volume datasets. Their method supports direct lighting, volumetric shadows and indirect lighting effects with interactively changing viewpoints, lighting and isovalues. Still, the complex scattering behavior inside the volume is ignored. Sharing the same motivation, Beason et al. [21] proposed an enhanced isosurface GI rendering method that further supports translucency and caustics, but this method is limited to static lighting. Ritschel [22] incorporated a GPU-based hierarchical visibility approximation into the precomputation to quickly compute the visibility inside the volume, but it only handles low-frequency effects. In comparison, our PVRT method can generally support various volume GI effects with few constraints for the input volume dataset. Furthermore, the precomputation-based routine has also been adopted in the offline rendering of the MS effect in hair [23].

**Real-time Volume GI Rendering** Kaplanyan et al. [24] developed a cascaded light propagation volume (LPV)



(a) Smoke Plume,  $300 \times 300 \times 150$ , 45fps (b) Bone,  $512 \times 512 \times 512$ , 40fps (c) Smoke Wall,  $400 \times 300 \times 100$ , 42fps (d) Machine Room,  $417 \times 345 \times 60$ , 47fps

Fig. 1. Examples of real-time rendering including volume GI effects using the PVRT method.

method to approximate indirect lighting very efficiently, which can further be extended to render SS effects in volumetric media. However, the LPV method only supports low-frequency effects in coarse resolution volumes, and it is non-trivial to support complex lighting with this method. For optically thick media, like smoke, the MS effect, which is much more costly compared with SS, dominates the overall appearance. Recently, Engelhardt et al. [25] suggested to create a set of virtual point lights (VPLs) within the medium and combine the SS contributions from all VPLs to render MS in real-time. Unfortunately, the limited number of VPLs usually leads to low-frequency effects. Several GPU-accelerated approaches [26], [27] have been proposed to render high-quality volumetric smoke under low-frequency environment light. For comparison, our PVRT method can render the MS effect under an all-frequency environment light in real-time.

Hernell et al. [28] presented an efficient method to compute ambient and emissive tissue illumination, which exploits ray casting to determine ambient illumination, and adopts a multi-resolution volume to speed up rendering. Similarly, to render simple scattering and ambient occlusion for isosurfaces, a GPU-based Monte Carlo ray casting method was proposed in [29]. When compared with these methods, our PVRT method is superior in its support for complex scattering, especially MS effects, for general volume data formats under complex lighting conditions. Kniss et al. [30] proposed a shading model to interactively render volumes, which generates volumetric shadows and forward scattering through translucent materials. Ropinski et al. [31] presented a physically more accurate optical model for volume rendering, which incorporates scattering, shadowing and specular reflections. However, applying these volume shading models to handle complex environmental lighting and general scattering effects is non-trivial. To render high-quality volume shadows, Kronander et al. [32] presented a method which encodes local and global volumetric visibility into an efficient multi-resolution grid over the extent of the volume to achieve real-time performance. Compared with it, our PVRT method can generate various volume GI effects, not just shadow, in real-time.




Fig. 2. The pipeline of the PVRT method for rendering multiple scattering.

$p$	A 3d point within medium
$p_{in}, p_{out}$	Entry/Exit point in medium
$x$	A 3d point
$\omega_i, \omega_o$	Incident, outgoing radiance direction
$S$	Sphere of directions
$D(p)$	Volume density field at $p$
$\kappa_s(p)$	Scattering coefficient
$\kappa_a(p)$	Absorption coefficient
$\kappa_t(p)$	Extinction coefficient
$\Omega(p)$	Albedo of participating medium
$\tau(p, x)$	Transmittance from $u$ to $x$
$\tau_\infty(p, \omega)$	Transmittance from infinity to $x$
$f(u, \omega_o, \omega_i)$	scattering phase function at $u$
$L_{in}(\omega)$	Incoming radiance from light source
$L_{out}(x, \omega)$	Outgoing radiance arriving at $x$
$L_{ri}(x, \omega)$	Reduced incident radiance
$L_m(x, \omega)$	Media radiance
$J(x, \omega)$	Scattered radiance
$L_{re}(\omega)$	Radiance exitance function
$T$	Radiant exitance matrix

Fig. 3. Nomenclature

### 3 RADIANCE TRANSPORT IN VOLUMETRIC MEDIA

Before the introduction of our PVRT method, we briefly review radiance transport in volumetric media. Following the definition in [5], the radiance quantities that we use are listed in the table of Fig.3. The input volume data is modeled as a sequence of volumetric density fields. Regular mesh grids are used to store the volumetric density field at cell centers. Considering a 3D position  $p$  in the volumetric media (shown in Fig.4), we represent the volume density field of  $p$  at each frame as  $D(p)$ , the absorption coefficient as  $\kappa_a(p)$ , the scattering coefficient as  $\kappa_s(p)$ , the extinction coefficient as  $\kappa_t(p) = \kappa_a(p) + \kappa_s(p)$ , and the

scattering albedo as  $\Omega(p) = \frac{\kappa_s(p)}{\kappa_t(p)}$ . Our method supports both environmental and local directional lighting. We use  $L_{in}$  to represent the incident radiance from the light source.

Fig.4 illustrates the radiance transport in volumetric media under environmental lighting. For view point  $x$ , the radiance reaching  $x$  along direction  $\omega_o$  is composed of two parts: the reduced incident radiance  $L_{ri}$  and the media radiance  $L_m$ :

$$L_{out}(x, \omega_o) = L_{ri}(x, \omega_o) + L_m(x, \omega_o) \quad (1)$$

The reduced incident radiance  $L_{ri}(x, \omega_o)$  describes the source illumination  $L_{in}(\omega_o)$  that arrives directly at  $x$  along direction  $\omega_o$  with some attenuation by the medium:

$$L_{ri}(x, \omega_o) = \tau_\infty(x, \omega_o)L_{in}(\omega_o) \quad (2)$$

where  $\tau_\infty(x, \omega_o)$  is the transmittance from infinite lighting to  $x$  along the direction  $\omega_o$ , that can be computed as  $\exp(-\int_x^\infty \kappa_t(u)D(u)du)$ . The media radiance




Fig. 4. Illustration of the radiance transport in volumetric media.

$L_m(x, \omega_o)$  represents the radiance energy that has scattered one or multiple times in the medium before arriving at  $x$  along the direction  $\omega_o$ . It is computed by integrating radiance contributions along the view ray:

$$L_m(x, \omega_o) = \int_{p_{in}}^{p_{out}} \tau(p, x)\kappa_t(p)D(p)J(p, \omega_o)dp \quad (3)$$

where source radiance  $J(p, \omega_o)$  represents the light that scatters at point  $p$  toward direction  $\omega_o$ . In usual non-emissive media, such as smoke, fog, etc., this source radiance consists of a single scattering (SS) term  $J_{ss}$  and a multiple scattering (MS) term  $J_{ms}$ :

$$J(p, \omega_o) = J_{ss}(p, \omega_o) + J_{ms}(p, \omega_o) \quad (4)$$

As shown in Fig.4, the SS term  $J_{ss}$  represents the reduced incident radiance  $L_{ri}$  whose first scattering interaction in the medium occurs at  $p$ :

$$J_{ss}(p, \omega_o) = \frac{\Omega(p)}{4\pi} \int_{\mathbb{S}} L_{ri}(p, \omega_i)f(p, \omega_o, \omega_i)d\omega_i \quad (5)$$

where  $f(p, \omega_o, \omega_i)$  is the phase function that defines the scattering distribution of the current media, and  $\mathbb{S}$  represents the set of different incident directions  $\omega_i$  that distributes on the local sphere centered at  $p$ . Same as Eq.2,  $L_{ri}(p, \omega_i)$  can be computed as  $\tau_\infty(p, \omega_i)L_{in}(\omega_i)$ . Notice that  $L_{ri}(x, \omega_o)$  can also be regarded as a special SS case. The  $L_{ri}(x, \omega_o)$

and  $J_{ss}(p, \omega_o)$  together contribute to the *direct illumination* effects for the volume.

The MS term  $J_{ms}$  represents the radiance that has been bounced more than one time before arriving at  $p$ , which accounts for the *global illumination* (GI) effects. In Fig.4, the red arrows represent the incident media radiance  $L_m$  contributed from MS:

$$J_{ms}(p, \omega_o) = \frac{\Omega(p)}{4\pi} \int_{\mathbb{S}} L_m(p, \omega_i)f(p, \omega_o, \omega_i)d\omega_i \quad (6)$$

Similarly to surface shading, the computation of GI effects is much more time-consuming when compared with direct illumination, so MS always becomes the bottleneck of the volume GI rendering.

---

#### Algorithm 1 Pseudocode of the PVRT method

---

- 1 Define final view radiance  $L_{out} = 0$
  - 2 **Precomputation:**
  - 3 Approximate the lighting using certain basis (SH or DSF).
  - 4 **FOR** each basis light:
  - 5     Part1: Direct illumination
  - 6     Sample  $N_s$  directions uniformly.
  - 7     **FOR** each voxel:
  - 8         Accumulate  $L_{ri}$  (Eq.2) from each direction.
  - 9         Compute radiant exitance vector.
  - 10     **ENDFOR**
  - 11     Part2: Global illumination
  - 12     Trace photons through volume grids.
  - 13     Splat absorbed photon's energy into neighbor voxels.
  - 14     **FOR** each voxel:
  - 15         Accumulate the photon's energy.
  - 16         Compute radiant exitance vector.
  - 17     **ENDFOR**
  - 18 **ENDFOR**
  - 19 Compress the radiant exitance matrices of all voxels using CPCA.
  - 20 **Rendering:**
  - 21 Project the current lightings into basis space.
  - 22 **FOR** each CPCA cluster:
  - 23     **FOR** each basis transfer matrix:
  - 24         Compute view-dependent radiance.
  - 25     **ENDFOR**
  - 26 **ENDFOR**
  - 27 Reconstruct and store radiance volume.
  - 28 GPU-based ray marching to evaluate  $L_{out}$ .
- 

## 4 ALGORITHM OVERVIEW

Like PRT, the PVRT method includes a precomputation stage and a rendering stage, as shown in Algorithm 1.

**Precomputation** The precomputation stage is involved with both direct and global illumination. To achieve efficient volume GI effects, one of our contributions is the seamless integration of the VPM method [2] into the precomputation stage, including photon traversal through the volume as well as the evaluation and compression of the radiance transfer matrices for each voxel. To achieve this goal, there exist three major challenges in the precomputation stage.

The first challenge is the selection of the basis functions for representing radiance from high-frequency incident lighting. As aforementioned, SH can only handle low-frequency lighting. To approximate higher frequency lighting, spherical wavelet basis functions (SWBFs) [17] or spherical radial basis functions (SRBFs) [4] are good



candidates. However, both have disadvantages as well. The SWBFs are not rotationally invariant, which prevents fast rotation of the basis functions. The SRBFs can be freely rotated but their construction and evaluation are quite slow because the SRBFs do not have the interpolation property and local support. The interpolation property means the coefficients are equal to the spherical function values at the sampling directions, so that reconstructing the spherical function can be rapidly achieved by interpolation. The basis functions with local support ensure that only a small number of coefficients are needed for evaluation of the spherical function at arbitrary directions. To fulfill these properties, we introduced a new discrete spherical function (DSF) that is rotationally invariant, interpolation-based and also has locally supported basis functions. The DSF can achieve similar accuracy to the SRBFs for approximating high-frequency environment light.

The second challenge is how to cast photons. Photon casting is simple if all the light sources and geometries are clearly defined. However, in the PRT pipeline, the lighting environment is approximated using a set of basis functions, such as SH and DSF. The values of these basis functions are distributed nonuniformly and the values can be either positive or negative. Therefore, we must design a general photon casting scheme for different types of basis functions. To generate photons with correct energy distributions, instead of casting photons from distant planes facing random directions with direction-dependent casting rate, we cast photons from the boundary surface of the volume and each photon has a random direction and also a direction-dependent energy value. This assures that all the generated photons will go into the volume and the energy distributions will converge to the spherical basis functions faster than a normal implementation. Such a photon casting strategy has been tested and worked well for both SH and DSF.

Another challenge is the efficient radiance transfer estimation using VPM. For classic PRT rendering, we can directly integrate the incident radiance on the surface over a hemisphere using backward ray tracing. Since photon mapping is a forward light tracing technique, a photon gathering step is required during the radiance estimation within the volume. If using an unmodified VPM that stores the photons inside a kd-tree, the computational cost for approximating the volume radiance transfer is extremely high. This is because the spherical numerical integrations for all pairs of irradiance and radiant exitance basis functions have to be considered. Moreover, sampling photons from the kd-tree is difficult to parallelize and memory efficiency is low. To address such a situation, we introduce a fast splatting-based radiance transfer estimation technique, wherein each photon's energy is splatted and accumulated into its neighbor voxels for further numerical integration.

**Rendering** To achieve real-time volume rendering, there are also some technical challenges. Volume rendering usually demands lots of samples along the rays, which is fast for simple scalar fields or low-dimensional vector fields. However, combined with the PRT framework, evaluating

each sample requires the processing of hundreds of variables. For example, if we use 6th order SH for irradiance and 4th order SH for radiant exitance, the size of the radiance transfer matrix is  $36 \times 16 = 576$  and the size of the light vector is 36. Then we need to perform a matrix-vector multiplication of this size for each sample and it is impossible to evaluate it in real-time for all voxels. Therefore, we introduce an early-evaluation technique to significantly reduce the rendering cost. It is achieved by using clustered PCA for transfer matrix compression and reducing the compressed basis transfer matrices into scalar values based on view and lighting changes before decomposition. The technical details are presented in Sec.7.

In general, as shown in Fig.2, the PVRT method approximates the lighting using a set of basis lights and each basis light corresponds to a basis function. During the precomputation stage, aided by the VPM technique, the radiance transfer estimation is performed to compute the radiant exitance for each voxel. To save storage and improve performance, we further compress the radiant exitance data for the whole volume. Thereafter, during the rendering stage, the precomputed radiant exitance data is efficiently decompressed and used to compute the final view-dependent radiance for each voxel. The PVRT method supports both environment light and directional light.

## 5 DISCRETE SPHERICAL FUNCTION

Before introducing the PVRT method, we first present the discrete spherical function (DSF) that has both the interpolation property and local support.




Fig. 5. Illustration of the discrete spherical function.

### 5.1 Definition

We sample a set of uniformly distributed directions  $\Omega = \omega_1, \omega_2, \dots, \omega_n$  over a unit sphere using icosahedral mesh subdivision, as shown in Fig.5. Using  $s_k$  to represent the sampled value at direction  $\omega_k$ , the DSF for an arbitrary direction  $\omega$  is defined as:

$$S(\omega) = \sum_{k=1}^n \frac{W_h(\omega, \omega_k)}{\sum_{i=1}^n W_h(\omega, \omega_i)} s_k = \sum_{k=1}^n \bar{w}_k(\omega) s_k \quad (7)$$

where  $\bar{w}_k(\omega)$  is the DSF basis function, which represents a normalized weight.  $W_h(\omega_i, \omega_j)$  is a symmetric kernel function, and it is defined as:

$$W_h(\omega_i, \omega_j) = \begin{cases} (h^2 - |\theta_{ij}|^2)^3, & |\theta_{ij}| < h \\ 0, & |\theta_{ij}| \geq h \end{cases} \quad (8)$$

Where  $\theta_{ij} = \cos^{-1}(\omega_i \cdot \omega_j)$  is the angle between any two directions  $\omega_i$  and  $\omega_j$ .  $h$  is defined as  $h = \min_{i,j} \theta_{k_i k_j}$ ,

which is the minimum angular distance between any pair of sampled directions  $\omega_{k_i}$  and  $\omega_{k_j}$  in  $\Omega$ . As shown in Fig.5,  $h$  can be regarded as the radius of the local impact region of each sampled direction. The definition of  $W_h$  follows the kernel definition in smoothed particle hydrodynamics [33], which is widely applied for fluid simulation. Compared with the Gaussian kernel, it has been proven to be more efficient for evaluation and has compact support in local region.

With such a definition, the kernel function  $W_h$  has the following properties:

- **Local Support**

$$W_h(\omega_i, \omega_j) \begin{cases} > 0, \theta_{ij} < h \\ = 0, \theta_{ij} \geq h \end{cases}$$

As shown in Fig.5,  $W_h$  provides compact support in local region.

- **Monotonicity**

$$\frac{dW_h(\omega_i, \omega_j)}{d\theta_{ij}} \leq 0$$

- **Interpolation** Based on the definition of  $h$ , if we pick up any two sampled directions  $\omega_{k_i}$  and  $\omega_{k_j}$  in  $\Omega$ , then

$$W_h(\omega_{k_i}, \omega_{k_j}) = \begin{cases} 1, k_i = k_j \\ 0, k_i \neq k_j \end{cases}$$


Therefore,  $W_h$  guarantees that when  $\omega = \omega_k$ ,  $S(\omega_k) = s_k$  and the DSF function can recover the sampled original value. When  $\omega$  is located between several samples, as shown in Fig.5, its value is the interpolation of its neighbor samples (inside red dot circle) weighted by  $W_h$ .

Eq.7 actually represents the reconstruction of the spherical function using DSF and  $s_k$  can be regarded as the corresponding coefficients for the DSF bases. For other bases, such as SH and SRBF, computing the coefficients is a process of projecting the spherical function onto basis space and such a projection step is usually time-consuming. However, for the DSF  $s_k$  is just a sampling of the spherical function, so the projection step is very efficient. Actually computing  $s_k$  is not a simple point sampling, and to balance the frequency between the spherical function and the DSF we introduce a prefiltering step for the spherical function which will be discussed in Sec.8. Relying on the local support property of DSF, the evaluation of Eq.7 for arbitrary directions can be efficiently computed because only a limited number of bases need to be considered. The definition of the DSF basis  $\bar{w}_k(\omega)$  is a normalized weight for local interpolation, so  $\bar{w}_k(\omega)$  is invariant during the rotation of the spherical function. Therefore, when rotating the environment light, we only need to rotate the prefiltered environment light and sample it again to get the new coefficient  $s_k$ , which is simple and efficient.

## 5.2 Lighting Approximation Using DSF

To approximate the lighting using this basis function, we project the original lighting function onto basis space and

compute the corresponding coefficient for each basis. With these coefficients, afterward we can reconstruct the value of the lighting function at arbitrary directions by evaluating the basis functions. Comparing the reconstructed value with the original value can help us verify the accuracy of the lighting approximation. We perform such a lighting reconstruction experiment [16] to test the DSF.



(a) Input EM (b) 162, 13.2% (c) 642, 8.9% (d) 2562, 7.3%

Fig. 6. Comparison of the approximation accuracy using different number DSF bases for a synthetic high-frequency EM (a). The 1st row lists the reconstruction results, and the 2nd row shows the approximation errors ('red' and 'blue' represent the positive and negative values, respectively.). The number of DSF bases and  $L^2$  errors are given for (b) (c) and (d).

The initial directional sampling of the DSF is based on icosahedral mesh generation, and the number of samples over the sphere  $\Omega$  can be 42, 162, 642, 2562... depending on the order of icosahedral subdivision. To select the optimal number of the DSF bases to achieve a good balance between approximation quality and performance, we reconstruct a synthetic high-frequency "checkerboard" environment map (EM) with different numbers of DSF bases and compare the results using standard  $L^2$  errors between the pixels in the original cube map and the pixels in the reconstructed one. As shown in Fig.6, the reconstruction quality essentially improves by increasing the number of DSF bases. Notice that using 2562 bases rather than 642 bases only improves quality by a factor of 1.6%, but multiplies the total number of basis lights by four. Therefore, we choose the optimal number of DSF bases to be 642.

To compare the approximation quality between the DSF and other existing bases, in Fig.7, we test the lighting reconstruction for both synthetic and real EM lights using different bases with a similar number of basis functions. Clearly, DSF can achieve the best approximation quality with lowest  $L^2$  error. To verify the high-frequency rendering effects using different bases, in Fig.8 we test the shadow effects under directional light. It is obvious that using a low-frequency SH basis misses the high-frequency shadow details and the ringing artifacts of SH darken the whole scene. Both DSF and SRBF achieve good high-frequency shadow effects.

The statistical results in our experiments show that




Fig. 7. Comparison of the approximation results for the synthetic (a) and real EMs (e,i,m). The 2nd column uses the 26th order SH with 676 coefficients, the 3rd column uses SRBF with 642 coefficients, and the 4th column uses DSF with 642 coefficients. The  $L^2$  errors are given for all the reconstructed results.




Fig. 8. Comparison of the shadow reconstruction accuracy using different bases.

dealing with an arbitrary EM light with size  $6 \times 256 \times 256$ , for each SH, SRBF and DSF basis, the average timing of the projection step is 539ms, 91ms, 2ms, respectively, and the average timing of evaluation for an arbitrary direction is 0.369ms, 0.059ms, 0.013ms, respectively. It is clear that SH is the most time-consuming basis for both projection and evaluation, and hence normally we should not use many SH bases to approximate high-frequency functions. SRBF computation involves numerical integration [4], and therefore its time cost is much higher than the DSF, especially

for the projection step. The DSF performs very efficiently for both steps because its interpolation property and local support ensure fast light projection and evaluation.

## 6 PRECOMPUTING RADIANT EXITANCE

In this section, we describe the precomputation stage of the PVRT method in detail. Since the environment light is more challenging, we use it for explaining the algorithmic steps.

### 6.1 Radiance and Lighting Approximation

As described in [3], a low-frequency spherical function can be approximated using  $s$ th order spherical harmonics

$$L(\omega_i) = \sum_{l=0}^{s-1} \sum_{t=-l}^l c_l^t Y_l^t(\omega_i) \quad (9)$$

where  $Y_l^t$  are the orthogonal real spherical harmonics (SH) basis functions and  $c_l^t$  are their corresponding coefficients. The radiant exitance of each voxel is actually a 2D spherical function that has a corresponding value for each direction over the unit sphere based on the phase function within the current voxel. Arbitrary phase functions are supported for computing radiant exitance in our method. However, storing the radiant exitance values can be challenging. Based on the fact that multiple scattering (MS) inside the volume exhibits low-frequency behavior [1], we choose the SH basis to approximate the 2D spherical function of the radiant exitance for each voxel. Depending on the current phase function, the selected order of SH varies. Following [26], we adopt 1st order SH for the isotropic phase function and 4 – 6th order SH for the anisotropic phase function.

The low-frequency environment light can also be well approximated using the low order SH (i.e.  $l \leq 6$ ) basis [3]. For the high-frequency environment light we use the aforementioned DSF for lighting approximation, and the radiance from arbitrary direction  $\omega_i$  can be represented as

$$L(\omega_i) = \sum_{k=1}^n \bar{w}_k(\omega_i) s_k \quad (10)$$

where  $\bar{w}_k(\omega_i)$  is the DSF basis function and  $s_k$  is the radiance value at the  $k$ th pre-sampled direction. After the approximation, the environment light is converted into a set of basis lights, and each basis light can be treated as distant spherical lighting.

### 6.2 Direct Illumination

For direct illumination, the incident energy for each voxel is just the reduced radiance  $L_{ri}$  (Eq.2), so photon tracing is unnecessary for its precomputation. For a voxel at position  $p$ , when using the SH or DSF basis to approximate incident lighting, the incident reduced radiance in  $\omega_i$  from each basis light can be computed as  $L_{ri}(p, \omega_i) = \tau_\infty(p, \omega_i) Y_l^m(\omega_i)$  or  $L_{ri}(p, \omega_i) = \tau_\infty(p, \omega_i) \bar{w}_k(\omega_i)$ . In the precomputation stage, for each basis light we uniformly sample  $N_s$  directions either over a unit sphere (for SH) or inside a



small solid angle around the current basis light (for DSF). Considering all the  $L_{ri}$ s from  $N_s$  directions, each  $L_{ri}$  is multiplied with the phase function value at the current voxel position to get the radiant exitance  $L_{re}$ . As mentioned above,  $L_{re}$ , which is a 2D spherical function, is non-trivial to store directly. Therefore, we use SH to approximate  $L_{re}$  with good accuracy. For each basis light, the radiant exitance of each voxel can thus be approximated and stored as a vector. After processing all basis lights, the radiant exitance of each voxel finally forms a *radiant exitance matrix*. Methods for determining the value of  $N_s$  will be discussed in Sec.8.

### 6.3 Global Illumination using VPM

In this section, we focus on how to address the second and third challenges that are mentioned in Sec.4 for the seamless incorporation of the VPM method into the PRT framework and the achievement of efficient precomputation for the rendering of volume GI effects.

#### 6.3.1 Volume Photon Sampling under Spherical Lighting

As is the case for the standard VPM method, the first step is to generate the photons for each basis light. Our basic idea is to cast photons at the boundary cells (voxels) of the volume with varying flux energy for each photon. The photon generation fulfills the following requirements: (1) the initial directions of all the generated photons in each boundary voxel are uniformly distributed on the unit sphere, and (2) each photon's energy is proportional to the radiance from the basis light in the photon's direction. As shown in




Fig. 9. Illustration of the volume photon sampling.

Fig.9, for each boundary voxel in the volume we generate  $k$  photons from the boundary surfaces of the voxel. Each photon starts from a random position over the boundary surface and has a random direction  $\omega$  that is uniformly distributed over the sphere. For each photon under SH basis light  $Y_l^t$ , the flux energy is:

$$\Phi(\omega) = Y_l^t(\omega) \Delta l^2 (\omega \cdot \mathbf{n}) d\omega \quad (11)$$

where  $\Delta l$  and  $\mathbf{n}$  are the side length and the outward normal of the boundary voxel, respectively. For each DSF basis light  $\bar{w}_k$ , the flux energy of each photon is:

$$\Phi(\omega) = \bar{w}_k(\omega) \Delta l^2 (\omega \cdot \mathbf{n}) d\omega \quad (12)$$

Some of the generated photons may immediately exit the volume region. Such photons are discarded before further computation. It is obvious that the number of generated photons increases as the resolution of volume data increases. To reduce the number of photons when dealing

with high-frequency lighting, rather than distributing photon's direction  $\omega$  over the whole unit sphere, we may restrict the  $\omega$  inside a small differential solid angle for each DSF basis light. Therefore, although the number of DSF bases is usually much higher than the number of SH bases, much fewer photons are actually needed for each DSF basis light.

#### 6.3.2 Volumetric Photon Tracing

Starting from the boundary voxel, the photons travel through the volume with a pre-defined step size. At each step, a photon can either pass unaffected through the medium inside current voxel, or be scattered or absorbed. We choose the side length of the voxel  $\Delta l$  as step size for photon propagation so that the maximum number of potential steps matches the volume resolution. The probability of an absorption or scattering event occurring at point  $p$  is integrated along the photon path from the entry point  $p_{in}$  to  $p$ :

$$P(p) = 1 - \exp \left\{ - \int_{p_{in}}^p (\sigma_a + \sigma_s) \rho(x) dx \right\} \quad (13)$$

where  $\rho$  is the density field function for volume media, which is often obtained from the original volume data through a user-defined mapping.  $\sigma_a$  and  $\sigma_s$  are absorption and scattering coefficients, respectively. When a random number  $0 \leq X \leq 1$  satisfies  $X \leq P$ , Russian Roulette is used to determine whether it is an absorption or scattering event based on the probability values  $\sigma_a/(\sigma_a + \sigma_s)$  and  $\sigma_s/(\sigma_a + \sigma_s)$ . The new scattering direction of a photon is chosen using importance sampling based on the phase function within the current voxel [2], and arbitrary phase functions can be supported in our method.

#### 6.3.3 Photon Splatting for Radiance Estimation

To avoid creating a hierarchical data structure to store photons [2], rather than gathering neighbor photons' energy for each voxel, the PVRT method chooses to splat each photon's energy into its neighbor voxels [14]. For an absorbed photon  $p$  with direction  $\omega$ , its energy  $\Phi_p(\omega)$  (Eq.11 or 12) is splatted into the photon's neighbor voxels weighted by the distance between the photon and voxel. Such a splatting process has an equivalent effect to the standard gathering process [2]. To ensure both efficient evaluation and compact support in local regions, again we adopt the polynomial kernel [33] as the weighting function:

$$W(d) = \begin{cases} \frac{315}{64\pi h^9} (h^2 - |d|^2)^3, & |d| < h \\ 0, & |d| \geq h \end{cases} \quad (14)$$

where  $d$  is the distance between the voxel and the photon  $p$ .  $h = c\Delta l$  is the kernel radius that is set based on the grid resolution, and  $c$  is a user-defined constant (e.g.  $c = 4$ ).

After the splatting, we compute the radiant exitance for each neighbor voxel based on its phase function  $f(x, \omega_o, \omega_i)$ . With the weighted photon energy  $W(d)\Phi_p(\omega)$  and photon's direction  $\omega$ , the radiant exitance of each neighbor voxel can be represented as  $L_{re}^p(\omega_o) = W(d)\Phi_p(\omega)f(x, \omega_o, \omega)$ . After processing all the photons, the final radiant exitance of each voxel is  $L_{re}(\omega_o) =$

$\sum_{p=0}^Q W(d)\Phi_p(\omega)f(x, \omega_o, \omega)$ , where  $Q$  is the number of the neighboring photons. Again, we use SH to approximate  $L_{re}(\omega_o)$  and the radiant exitance of each voxel is approximated and stored as a vector. With the radiant exitance for each voxel, we can evaluate the incident radiance for arbitrary view rays in the subsequent volume ray casting process.

#### 6.4 The Radiant Exitance Matrix and Its Compression

For each voxel, the radiant exitance under each basis light is represented as a vector with length  $n$ . After processing  $m$  basis lights, the radiant exitance from all the basis lights forms a  $m \times n$  matrix  $T$ , called the *radiant exitance matrix*, which can be regarded as a linear transformation that transforms the input lighting vector  $L_{in}$  into final radiant exitance vector  $I_e$ :

$$I_e = TL_{in} \quad (15)$$

Here the length of  $L_{in}$  and  $I_e$  is  $m$  and  $n$ , respectively, which represent the number of used bases. Computing  $T$  for each voxel is the major task during the precomputation stage. If using a 6th order SH for the environment light and a 4th order SH for the anisotropic phase function (i.e.  $m = 36$  and  $n = 16$ ), the dimension of  $T$  will be  $36 \times 16$ . Storing a  $36 \times 16$  matrix for each voxel is impractical, so we need to further compress the precomputed data.

We adopt the clustered principal component analysis (CPCA) method [34] to achieve further compression. All the matrices of the voxels are first clustered in signal space using vector quantization. Note that empty voxels with zero density gradients are excluded from the clustering and zero weights are assigned directly. The total number of clusters are adjustable by compromising between quality and speed, and a relatively optimal cluster number is 128 based on our experiments (Sec.9). Thereafter, within each cluster, the standard principal component analysis (PCA) is performed to compute the mean matrix, the set of the basis matrices and the set of corresponding coefficients for each voxel's radiant exitance matrix. Using CPCA, the transfer matrix  $T$  at a certain voxel  $t$  can be approximated through

$$T \approx T_{mean}^k + \sum_{i=1}^N w_i^t T_i^k \quad (16)$$

where  $k$  is the index of the cluster that  $t$  belongs to,  $T_{mean}$  is the mean matrix of the  $k$ th cluster,  $T_i^k$  is the  $i$ th PCA basis matrix, and  $w_i^t$  is the corresponding coefficient.  $N$  is the number of PCA basis transfer matrices, and our experiments demonstrate that  $N = 4$  works well in most cases (Sec.9). Putting Eq.16 into Eq.15, the linear radiance transformation at each voxel is approximated as:

$$I_e \approx (T_{mean}^k + \sum_{i=1}^N w_i^t T_i^k) L_{in}. \quad (17)$$

After the CPCA compression, the storage of the precomputation results of each voxel includes only one integer

cluster index  $k$  and  $N = 4$  PCA floating point coefficients. For each matrix cluster, we also store the mean matrix  $T_{mean}$  and  $N = 4$  basis matrices. Overall, such a strategy saves lots of storage and also ensures real-time performance for further radiance estimation in the ray casting step. The accuracy of our compression strategy is discussed in Sec.9.

## 7 REAL-TIME VOLUME RENDERING

The major steps in the real-time rendering stage consist of (1) view-dependent volume radiance estimation, and (2) volume ray casting. All the operations in the rendering stage can be efficiently performed on the GPU.

### 7.1 Lighting Approximation

Incident lighting is projected into basis function space, either SH or DSF, and can be represented as a vector of projection coefficients  $L_{in}$ . This is a standard step in PRT-based methods and is performed in each frame.

### 7.2 View-dependent Volume Radiance Estimation

After updating the lighting environment, in each cluster  $k$ , we multiply the new light coefficient vector  $L_{in}$  to the mean matrix  $T_{mean}^k$  as well as all the basis transfer matrices  $T_i^k$ .

$$I_{mean}^k = T_{mean}^k L_{in} \quad I_i^k = \sum_{i=1}^N T_i^k L_{in} \quad (18)$$

where  $I_{mean}^k$  and  $I_i^k$  represent the transferred mean and basis radiant exitance vector of the  $k$ th cluster, respectively. The final radiance vector  $I_e$  at voxel  $t$  can be recovered by:

$$I_e = I_{mean}^k + \sum_{i=1}^N w_i^t I_i^k \quad (19)$$

where  $w_i^t$  is the PCA coefficients for voxel  $t$ . The straightforward strategy is to store the  $I_e$  for each voxel for further radiance estimation. However,  $I_e$  is a vector of length  $n$  and directly saving it for all voxels requires  $n \times Size_{volume}$  GPU memory, which is costly. Furthermore, evaluating Eq.19 is also inefficient for all the voxels in each frame. Notice that in each frame the current view info is known before volume ray casting. As shown in Line 22 – 26 in Algorithm 1, for the  $k$ th cluster we can first compute the view-dependent radiance value  $L_V^k$  for  $I_{mean}^k$  and each  $I_i^k$ , respectively, by evaluating the basis functions with the view direction. When using the SH basis,  $I_V^k$  is computed as (using Eq.9):

$$L_V^k = \sum_{l=0}^{s-1} \sum_{t=-l}^l I_{l(l+1)+t+1} Y_l^t(\omega_{x \rightarrow x_V}) \quad (20)$$

when using DSF basis, it is computed as (using Eq.10):

$$L_V^k = \sum_{k=1}^n \bar{w}_k(\omega_{x \rightarrow x_V}) s_k \quad (21)$$

where  $x \rightarrow x_V$  is a unit vector from voxel position  $x$  to current view point  $x_V$ . Thereafter, for each voxel, its view-dependent radiance value  $L_V$  can be efficiently computed



as the linear combination of the  $L_V^k$ s weighted by the PCA coefficients of current voxel. With such an early-evaluation strategy, we only need to save one floating point value for each voxel, and the storage for all the voxels is reduced to be just one additional *radiance volume* texture.

### 7.3 Volume Ray Casting

The aforementioned radiance estimation for each voxel leads to a very simple ray casting stage. The traditional GPU-based volume ray casting method [35] can be applied to render the volume data. In addition, we just need to sample the radiance  $L_V$  from the *radiance volume* texture at each sampling point along the ray.

## 8 IMPLEMENTATION

In this section, we will discuss several issues and the corresponding solution strategies we use for implementation.

### 8.1 Environment Map Prefiltering for DSF

We adopt 642 DSF bases to approximate the high-frequency lighting. However, the size of input EM usually contains more than 642 pixels. For example, a standard EM might have  $6 \times 256 \times 256$  pixels. Sampling 642 values discretely from the EM will probably miss important lighting energy, especially for high-frequency lighting. Therefore, to match the frequency between the EM and DSF, we introduce a prefiltering step for the EM. The directions of DSF bases are uniformly distributed over the sphere (Fig.5), and since each EM pixel corresponds to a direction, we can map the DSF basis direction to the pixels in EM. To ensure the prefiltering results cover the whole spherical domain, we choose a spherical Gaussian with bandwidth equal to minimum angular distance  $h$  (defined in Eq.8). To perform the prefiltering, we use an image space Gaussian centered at the pixel corresponding to the DSF basis direction with the bandwidth computed by mapping  $h$  into image space. The prefiltering step can be computed in real-time. After the prefiltering, each pixel value in the prefiltered EM is a coefficient for a DSF basis. The prefiltering step only needs to be performed once unless the content of input EM changes (e.g. video EM). When rotating the EM lighting, we just need to rotate the prefiltered EM and then sample it directly to get the new coefficients after rotation.

### 8.2 Directional Light

It is well known that the directional light can be easily represented by the SH bases evaluated at the light direction. Approximating the local directional light using DSF, we only need to find the three nearest DSF directions for the current light direction based on angular distance, and compute three corresponding weights using barycentric interpolation. Then, the directional light can be approximated by the linear combination of the three DSF basis lights.

### 8.3 Number of Photons for Each Voxel

In the PVRT method, the photons are emitted from boundary voxels and the number of photons per voxel  $N_P$  is controlled by the user.  $N_P$  affects the rendering quality and it scales well as the resolution of the input volume dataset changes. For example, if we set  $N_P$  to be 100, a  $100^3$  volume would cast  $6 \times 100^3 = 6$  million photons in total because there are only  $6 \times 100^2$  boundary voxels. As discussed in Sec.6.3.1, when using DSF  $N_P$  is much smaller compared with using SH. Based on our experimental results, we found that  $N_P = 10$  and  $N_P = 100$  achieve a good balance between quality and speed for DSF and SH, respectively. More results and discussions will be shown in Sec. 9.

### 8.4 Performance Optimizations

There are several performance optimization strategies: (1) When computing direct illumination during precomputation, each voxel marches  $N_s = 642$  rays through the volume to gather incident lighting. If the volume resolution is  $512^3$ , the total number of ray samplings will be 86 billion. To accelerate this process, we skip all the empty voxels ( $D(p) \leq 0$ ). Also, a ray is terminated early if a high cumulative density is reached during its marching procedure. (2) We use pre-calculated lookup tables to store SH basis function values instead of evaluating them at runtime, which greatly accelerates the photon splatting operation and radiance integration. (3) In the precomputation stage, the computation for different basis lights and different photons is highly independent. Therefore, it is easy to parallelize the computation of each individual case (i.e. certain basis light) using the evolving many-core hardware. Our CUDA-based parallel implementation for precomputation achieves more than two orders of magnitude speed improvement over the single-threaded CPU implementation.

### 8.5 System Integration

It is easy to integrate our method into existing rendering frameworks. Our renderer consists of two modules including a preprocessor and a real-time renderer. The preprocessor is called once, which takes the density volume together with several parameters as input and outputs the volume radiance transfer data. Then the real-time renderer takes the volume data, the volume radiance transfer data as well as other rendering parameters as input and outputs the final image. In our implementation, we fit our renderer into an OpenGL-based application framework and created several graphics user interfaces for adjusting the parameters.

## 9 RESULTS AND DISCUSSIONS

In this section, we report on the accuracy and performance of the PVRT method. Further, we compare the results of our method with the references generated by the VPM [2]. Our method is implemented in OpenGL and CUDA [36]. All results were rendered on a laptop equipped with an Intel® i7 740QM 1.73GHz CPU, an NVIDIA® Geforce GTX 460M GPU and 6GB of physical memory. All the results,

if not specifically mentioned, utilize DSF with 642 bases to approximate incident lighting. The number of CPCA clusters is 128, and the number of PCA bases is 4.

TABLE 1  
Performance data.

VolData	DataRes	NumP (mil.)	Prep (sec.)	RV (ms)	Render (ms)
Smoke Plume	$300 \times 300 \times 150$	3.6	110	0.8	20
Smoke Wall	$400 \times 300 \times 100$	3.8	165	0.8	23
Aneurism	$256 \times 256 \times 256$	3.9	81	0.9	19
Bonsai	$256 \times 256 \times 256$	3.9	113	0.9	20
Bone	$512 \times 512 \times 512$	7.9	241	1	23
Part	$504 \times 504 \times 228$	9.7	395	1	25
Machine Room	$417 \times 345 \times 60$	3	73	0.7	21

The performance timings for different datasets are listed in Table 1. As shown, the number of sampled photons (col. *NumP*) increases as the resolution of volume dataset increases. The precomputation time (col. *Prep*) is related to both the number of photons and the density field occupation inside the volume that determines the effective voxels during ray/photon marching. Usually, more photons demand more precomputation time. Note, the precomputation time listed in Table 1 includes the timings of both single (SS) and multiple scattering (MS). The PVRT method aims at handling high-resolution volumes, and even with a  $504 \times 504 \times 228$  volume, the precomputation time is just around 6.5 minutes. Also note that the current test GPU is designed for laptops, and thus has much weaker performance when compared with the PC version. With state-of-the-art GPUs, our precomputation performance is expected to be further improved by at least a factor of 3. For the rendering stage, the step of generating the radiance volume texture is very efficient, and for all datasets we tested it only takes around 1 ms (col. *RV*). Relying on the radiance volume texture, the following ray marching can be performed in real-time (col. *Render*).

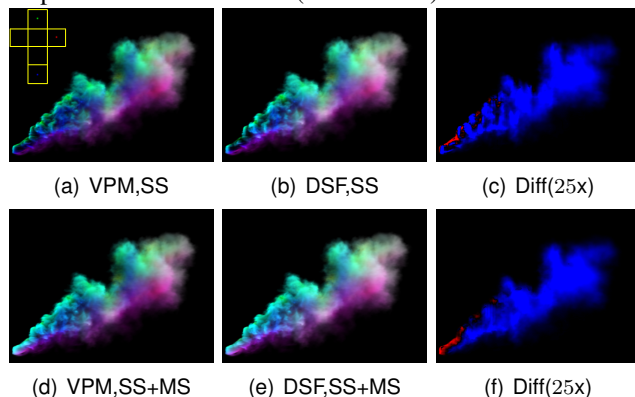


Fig. 10. Comparison between the ground truth generated by VPM and our results using DSF.

TABLE 2  
Statistics of different number of CPCA clusters.

No. Cluster	$L^2$	CPCA Time (s)
32	2.2%	12
64	1.3%	19
128	0.9%	31
256	0.8%	52

To validate the accuracy of our method, Fig.10 gives a comparison between our results using DSF and the references generated by VPM, which is treated as ground

truth. The first row shows the results with only direct illumination from SS, and the second row shows the full GI with both SS and MS. Visually, we can clearly distinguish the difference between with and without volumetric GI effects. The difference images show that our results are very close to the references. To compare the accuracy of different basis functions, Fig.11 shows the comparison between the references and the results generated by the PVRT method using different bases. In the first row, we utilized a synthetic environment light with three spherical lights in different colors, and in the second row we adopted a real EM. It is easy to notice that DSF reaches the best visual quality compared with other bases, especially SH. The  $L^2$  errors also numerically demonstrate this conclusion. To verify the




Fig. 12. Comparison of volume GI rendering with different volume resolution.

quality loss of using CPCA compression, we compared the rendering of a smoke dataset, wherein density distribution is random, using different numbers of CPCA clusters. The statistical data is shown in Table 2, which shows that using 128 clusters only introduces 0.1% more  $L^2$  error compared with using 256 clusters, while the CPCA compression using 256 clusters spends almost double the time compared with using 128 clusters. Hence, using 128 clusters for CPCA compression is optimal to achieve a good balance between performance and quality. To test the impact of

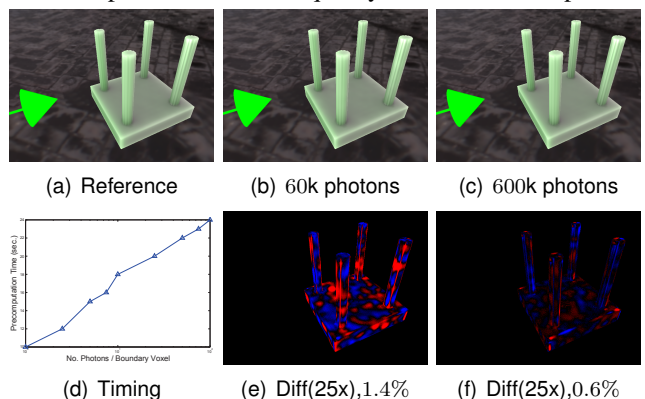


Fig. 13. Comparison of using different number of photons.

different input volume resolutions on the rendering quality, we spatially down-sampled the machine room dataset to half the resolution, and rendered it using directional light. Since the resolution decreases, the total number of photons and the precomputation time is reduced to be 0.75 mil. and 21 sec., respectively, compared with the original dataset (Table 1). Yet, as shown in Fig.12, the rendering quality also declines and the  $L^2$  error from the result using high-resolution volume is 4.3%. As mentioned in Sec.8.3, when using DSF, each boundary voxel shoots  $N_P = 10$  photons for each basis light. To verify the quality, in Fig.13, we

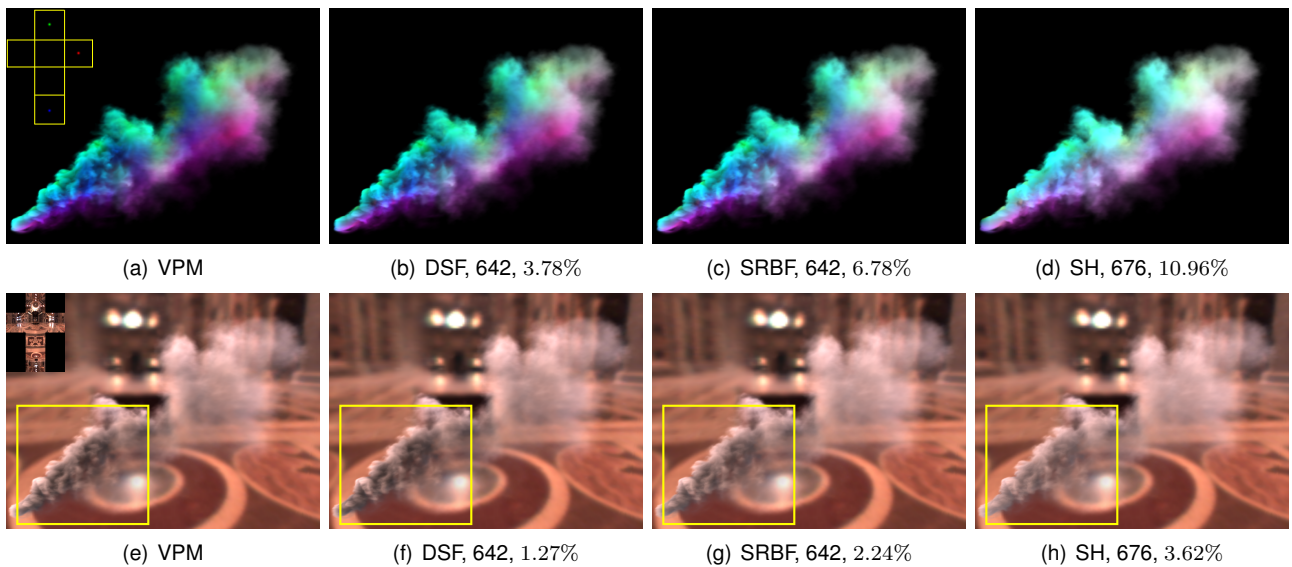


Fig. 11. Comparison between the references generated by VPM and the results using different bases.

compared the rendering results using  $N_P = 1$  and  $N_P = 10$  with the reference. This resolution of this Pillars volume is  $100^3$ , and hence the total number of boundary voxels is  $6 \times 100^2$ . Since the VPM strategy only contributes to the indirect illumination of the final rendering, visually it is not easy to distinguish the differences between using different numbers of photons. However, the difference images (10x) and the  $L^2$  errors clearly show the accuracy difference. Since when using  $N_P = 10$  photons, the  $L^2$  error is below 1%, and hence it provides sufficient visual quality at good performance. Further, Fig.13 (d) illustrates the trend of the precomputation time as a function of number of photons for each boundary voxel, which yields almost a linear curve.

We applied the PVRT method for various visualization scenarios. Fig.14 shows three study cases where our renderer is used for volume visualization. The first row shows the rendering of a bone volume dataset using conventional volume ray casting with local shading and our method (using DSF) with volume GI effects. In the first image, it is obvious that certain features, such as the spatial structure of the spine, cannot be clearly perceived under direct illumination. However, these features are correctly presented using our method with volume GI. In addition, by using both key light and back light, MS effects from the back light offer better perception of the structure within the spine. The second and third dataset is a CT scan of a bonsai tree and an aneurysm, respectively. It is noticeable that our result looks more 3D and provides a much better visual perception of the depth order of complex leaves or vessel structures. To demonstrate whether the high-frequency lighting is meaningful for visualization, in Fig.15, we tested the part dataset under environment light that is approximated using 6th SH and DSF with 642 bases, respectively. From this test, we can easily see that the rendering using low-frequency lighting smoothes out a great deal of the spatial features over the volume (shown in the yellow square), while the volume rendering results with high-frequency lighting preserves the important volume

features and hence provides a much better perception. Therefore, along with the real-time performance, the PVRT method offers high potential for improving the scientific and medical visualization quality.

**Limitations** In the precomputation stage of the PVRT method, the physical properties of each voxel, such as density field value, phase function, etc, can impact the transport of each photon and determine the final radiant exitance. Therefore, although we can render the volumetric effects in real-time, the physical properties of each voxel cannot be changed during rendering because the radiance transportation is precomputed. Also the indirect light transportation does not change according to the change of opacity mapping. However, since photon tracing is pursued for each basis light, the basis coefficients for current lighting can be freely changed. Hence we can support dynamic lighting, view and color transfer function change, in real-time (as shown in the accompanying video).

## 10 CONCLUSION AND FUTURE WORKS

In this paper, we present a precomputed volume radiance transfer (PVRT) method that is feasible to incorporate global illumination at real-time rates for rendering volume data. The added realism and clarity can greatly benefit practical visualization tasks. To simulate time-consuming multiple scattering, the PVRT method incorporates the volumetric photon mapping technique into the precomputation stage, and introduces a novel photon sampling strategy and photon splatting routine to compute the radiant exitance for each voxel. Our PVRT framework currently supports both distant spherical lighting and local directional lighting. To approximate all-frequency lighting, we further present the discrete spherical functions (DSF), which has a local interpolation property and is accurate and efficient for reconstruction. The PVRT method is practical for rendering high quality volume GI effects for various high-resolution volume datasets. Our experimental results demonstrate its

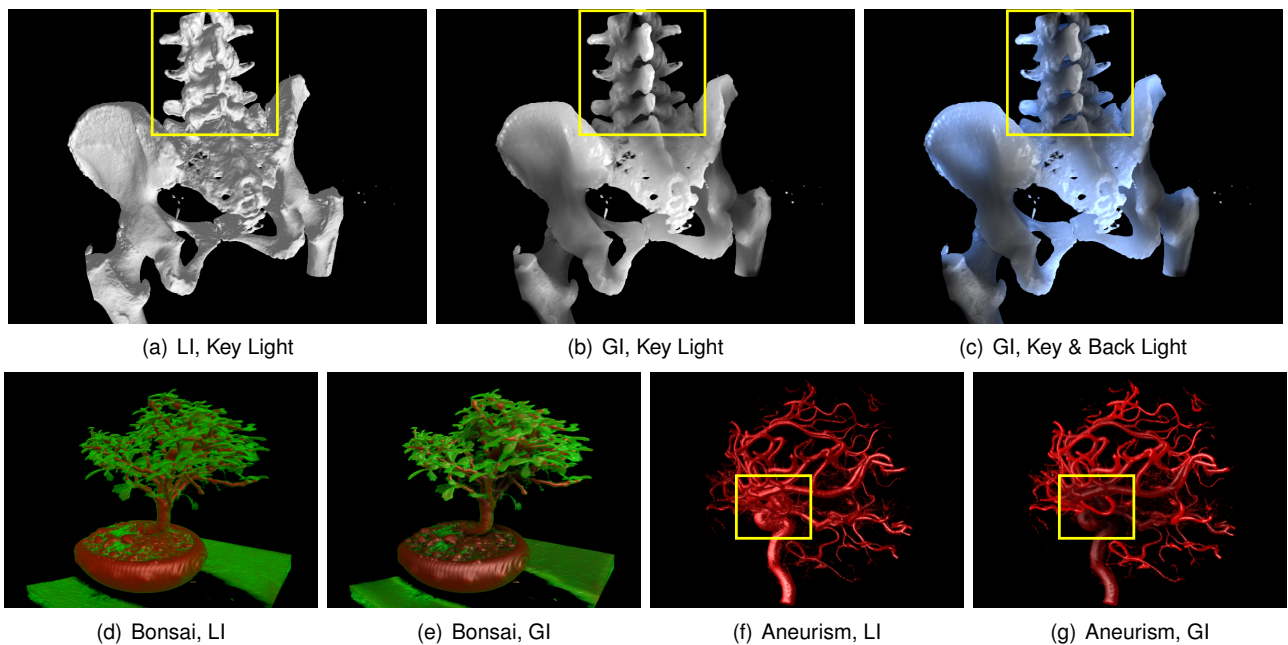


Fig. 14. Comparison of volume rendering results between local illumination (LI) and global illumination (GI) using our method.




Fig. 15. Comparison of volume GI rendering for visualization between using low-frequency SH (left column) and using high-frequency DSF (right column).

merits and show that it can benefit both interactive graphics and scientific visualization applications.

For future work, we aim at removing the current constraint of fixing the volume physical properties in pre-computation so that the user can interactively edit the volume dataset. Further, since the PVRT method inherits the routine of VPM, it inherently could support the rendering of volume caustics. As a next step, we would thus like to investigate high-quality volume caustics rendering using the PVRT method.

## ACKNOWLEDGMENTS

This research was sponsored in part by the U.S. National Science Foundation through grants OCI 0905008, OCI 0850566, and OCI 0749227, and also by the U.S. Department of Energy through grants DE-FC02-06ER25777 and


DE-FC02-12ER26072, program manager Lucy Nowell. We would like to thank volvis.org for the volume data and Paul Debevec for the environment maps. Finally, we sincerely thank the reviewers and Bob Miller for helping improve the manuscript.

## REFERENCES


- [1] J. T. Kajiya and B. P. Von Herzen, "Ray tracing volume densities," *SIGGRAPH Comput. Graph.*, vol. 18, pp. 165–174, 1984.
- [2] H. W. Jensen and P. H. Christensen, "Efficient simulation of light transport in scenes with participating media using photon maps," in *Proc. of SIGGRAPH*. ACM, 1998, pp. 311–320.
- [3] P.-P. Sloan, J. Kautz, and J. Snyder, "Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments," *ACM Trans. Graph.*, vol. 21, pp. 527–536, 2002.
- [4] Y.-T. Tsai and Z.-C. Shih, "All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation," *ACM Trans. Graph.*, vol. 25, pp. 967–976, 2006.



- [5] E. Cerezo, F. Perez-Cazorla, X. Pueyo, F. Seron, and F. Sillion, "A survey on participating media rendering techniques," *the Visual Computer*, 2005.
- [6] M. Hadwiger, P. Ljung, C. R. Salama, and T. Ropinski, "Advanced illumination techniques for gpu volume raycasting," in *ACM SIGGRAPH ASIA 2008 courses*, 2008, pp. 1:1–1:166.
- [7] J. Kautz, P.-P. Sloan, and J. Lehtinen, "Precomputed radiance transfer: theory and practice," in *ACM SIGGRAPH Courses*. ACM, 2005.
- [8] R. Ramamoorthi, "Precomputation-based rendering," *Foundations and Trends in Computer Graphics and Vision*, vol. 3, no. 4, pp. 281–369, 2009.
- [9] S. Chandrasekhar, *Radiative Transfer*. New York: Dover Publications, 1960.
- [10] M. Levoy, "Efficient ray tracing of volume data," *ACM Trans. Graph.*, vol. 9, pp. 245–261, 1990.
- [11] E. P. Laforune and Y. D. Willems, "Rendering participating media with bidirectional path tracing," in *Proc. of EGWR*, 1996, pp. 91–100.
- [12] H. E. Rushmeier and K. E. Torrance, "The zonal method for calculating light intensities in the presence of a participating medium," *SIGGRAPH Comput. Graph.*, vol. 21, pp. 293–302, 1987.
- [13] J. Stam, "Multiple scattering as a diffusion process," in *Proc. of EGWR*, 1995, pp. 41–50.
- [14] A. Boudet, P. Pitot, D. Pratumarty, and M. Paulin, "Photon splatting for participating media," in *Proc. of GRAPHITE*, 2005, pp. 197–204.
- [15] K. Zhou, Q. Hou, R. Wang, and B. Guo, "Real-time kd-tree construction on graphics hardware," *ACM Trans. Graph.*, vol. 27, pp. 126:1–126:11, 2008.
- [16] R. Ng, R. Ramamoorthi, and P. Hanrahan, "All-frequency shadows using non-linear wavelet lighting approximation," *ACM Trans. Graph.*, vol. 22, pp. 376–381, 2003.
- [17] W.-C. Ma, C.-T. Hsiao, K.-Y. Lee, Y.-Y. Chuang, and B.-Y. Chen, "Real-time triple product relighting using spherical local-frame parameterization," *Vis. Comput.*, vol. 22, no. 9, pp. 682–692, Sep. 2006.
- [18] P. Green, J. Kautz, W. Matusik, and F. Durand, "View-dependent precomputed light transport using nonlinear gaussian function approximations," in *Proc. of I3D*, 2006, pp. 7–14.
- [19] K. Zhou, Y. Hu, S. Lin, B. Guo, and H.-Y. Shum, "Precomputed shadow fields for dynamic scenes," *ACM Trans. Graph.*, vol. 24, no. 3, pp. 1196–1201, Jul. 2005.
- [20] C. Wyman, S. Parker, P. Shirley, and C. Hansen, "Interactive display of isosurfaces with global illumination," *IEEE Trans. Vis. Comp. Graph.*, vol. 12, no. 2, pp. 186–196, Mar. 2006.
- [21] K. M. Beason, J. Grant, D. C. Banks, B. Futch, and M. Y. Hussaini, "Pre-computed illumination for isosurfaces," in *Proc. of Visualization and Data Analysis*, 2006, pp. 98–108.
- [22] T. Ritschel, "Fast gpu-based visibility computation for natural illumination of volume data sets," in *Proc. of Eurographics*, 17-20, Ed., 2007.
- [23] J. T. Moon, B. Walter, and S. Marschner, "Efficient multiple scattering in hair using spherical harmonics," *ACM Trans. Graph.*, vol. 27, pp. 31:1–31:7, 2008.
- [24] A. Kaplanyan and C. Dachsbacher, "Cascaded light propagation volumes for real-time indirect illumination," in *Proc. of I3D*. ACM, 2010, pp. 99–107.
- [25] T. Engelhardt, J. Novak, and C. Dachsbacher, "Instant multiple scattering for interactive rendering of heterogeneous participating media," Karlsruhe Institut of Technology, Tech. Rep., December 2010.
- [26] K. Zhou, Z. Ren, S. Lin, H. Bao, B. Guo, and H.-Y. Shum, "Real-time smoke rendering using compensated ray marching," *ACM Trans. Graph.*, vol. 27, pp. 36:1–36:12, 2008.
- [27] Z. Ren, K. Zhou, S. Lin, and B. Guo, "Gradient-based interpolation and sampling for real-time rendering of inhomogeneous, single-scattering media," *Computer Graphics Forum*, pp. 1945–1953, 2008.
- [28] F. HERNELL, P. Ljung, and A. Ynnerman, "Efficient Ambient and Emissive Tissue Illumination using Local Occlusion in Multiresolution Volume Rendering," in *IEEE/EG Volume Graphics*, 2007.
- [29] C. R. Salama, "Gpu-based monte-carlo volume raycasting," in *Proc. of Pacific Graphics*, 2007, pp. 411–414.
- [30] J. Kniss, S. Premoze, C. Hansen, P. Shirley, and A. McPherson, "A model for volume lighting and modeling," *IEEE Trans. Vis. Comp. Graph.*, vol. 9, no. 2, pp. 150–162, Apr. 2003.
- [31] T. Ropinski, C. Döring, and C. Rezk-Salama, "Interactive volumetric lighting simulating scattering and shadowing," in *Proc. of PacificVis*, 2010, pp. 169–176.
- [32] J. Kronander, D. Jönsson, J. Low, P. Ljung, A. Ynnerman, and J. Unger, "Efficient visibility encoding for dynamic illumination in direct volume rendering," *To Appear in IEEE Trans. Vis. Comput. Graph.*, 2011.
- [33] M. Müller, D. Charypar, and M. Gross, "Particle-based fluid simulation for interactive applications," in *Proc. of SCA'03*, 2003, pp. 154–159.
- [34] P.-P. Sloan, J. Hall, J. Hart, and J. Snyder, "Clustered principal components for precomputed radiance transfer," *ACM Trans. Graph.*, vol. 22, no. 3, pp. 382–391, 2003.
- [35] J. Kruger and R. Westermann, "Acceleration techniques for gpu-based volume rendering," in *Proc. of IEEE Visualization 2003 (VIS'03)*, 2003, pp. 287–292.
- [36] NVIDIA, "Cuda programming guide," 2011, <http://developer.nvidia.com/cuda-downloads>.



**Yubo Zhang** is a Ph.D student of computer science at the University of California, Davis. He received his M.Phil. degree in applied mathematics from Hong Kong Baptist University in 2009 and B.S. degree in computational mathematics from Zhejiang University in 2007. His current research interests include scientific visualization, volume rendering and physically-based animation.



**Zhao Dong** is a postdoc at Program of Computer Graphics of Cornell University. He received his Ph.D. degree in Computer Graphics from Max-Planck-Institut Informatik, Germany in 2011, and M.S. and B.S degrees from Zhejiang University, China, in 2005 and 2001, respectively. His current research interests include real-time global illumination rendering, physically-based material modeling and rendering, and volume graphics.



**Kwan-Liu Ma** an IEEE Fellow, is a professor of computer science and the chair of the Graduate Group in Computer Science (GGCS) at the University of California, Davis. He leads the ViDi research group and directs the UC Davis Center for Visualization. Professor Ma received his PhD degree in computer science from the University of Utah in 1993. He was a recipient of the PECASE award in 2000. His research interests include visualization, high-performance computing, computer graphics, and user interface design. Professor Ma was a paper chair of the IEEE Visualization Conference in 2008 and 2009, and an associate editor of IEEE TVCG (2007-2011). He is a founding member of the IEEE Pacific Visualization Symposium and the IEEE Symposium on Large Data Analysis and Visualization. Professor Ma presently serves on the editorial boards of the IEEE CG&A, the Journal of Computational Science and Discoveries, and the Journal of Visualization. Contact him via email: [ma@cs.ucdavis.edu](mailto:ma@cs.ucdavis.edu).